vmware PRESS

Operationalizing VMware vSAN™

Kevin Lees, Chief Technologist for IT Operations Transformation, VMware

Paul Wiggett, Senior Technical Operations Architect, VMware

Foreword by **Duncan Epping**, Chief Technologist, VMware - Storage & Availability

vmware® PRESS

Operationalizing VMware vSAN™

Kevin Lees, Chief Technologist for IT Operations Transformation, VMware

Paul Wiggett, Senior Technical Operations Architect, VMware

Foreword by **Duncan Epping**, Chief Technologist, VMware - Storage & Availability

Operationalizing VMware vSAN™

Kevin Lees, Chief Technologist for
IT Operations Transformation, VMware

Paul Wiggett, Senior Technical
Operations Architect, VMware

Foreword by **Duncan Epping**, Chief Technologist, VMware - Storage & Availability

vmware® PRESS

VMWARE PRESS

Program Manager

Michael Haag

Technical Writer

William Rogers

Graphics Manager

Sappington

Production Manager

Sappington

Warning & Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an "as is" basis. The authors, VMware Press, VMware, and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

The opinions expressed in this book belong to the author and are not necessarily those of VMware.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA
Tel 877-486-9273 Fax 650-427-5001 www.vmware.com.

Table of Contents

List of Figures

List of Tables

About the Authors

Kevin Lees is the field Chief Technologist for IT Operations Transformation at VMware, focused on how customers optimize the way they operate VMware-supported environments and solutions. He is responsible for defining, communicating, and evangelizing VMware's IT Operations Transformation vision and strategy as it relates to operational (integrated organization, people, process, and application of VMware technology) approaches and best practices. Kevin also serves as an advisor to global customer senior executives for their IT operations transformation initiatives. Additionally, he leads the IT Transformation activities in VMware's Global Field Office of the CTO.

Kevin is an industry-recognized IT Operations Transformation thought leader, strategist, and evangelist. He was the global services delivery team lead for VMware's Cloud Practice from 2009 through mid-2011 and led many of VMware's early cloud implementations with Fortune 100 customers. Since 2011, Kevin has focused on initiating, building, and supporting IT Operations Transformation as a consultative delivery practice within VMware. Kevin has a combined 35+ years of IT system architecture, design, and integration, IT Operations, and IT consulting experience.

Kevin Lees can be reached at klees@vmware.com

Paul Wiggett is a Senior Technical Operations Architect working as part of the Transformation Services Practice at VMware. He works as a trusted advisor to some of VMware's largest customers across Europe, the Middle East and Africa assisting and guiding them in their digital transformation strategy and the subsequent transformation of their IT operations teams to a service-oriented and software defined approach.

Paul has over 15 years' experience in the IT industry and has a vast wealth of experience delivering large-scale cloud and infrastructure projects and helping deliver cutting edge solutions to customers across the globe. He is always looking to evangelize best practices and lessons learned and has presented at numerous conferences and C-level events both nationally and internationally throughout his career.

Paul Wiggett can be reached at
pwiggett@vmware.com

Acknowledgements

It takes the knowledge and resources of multiple individuals to create a book like *Operationalizing VMware vSAN®* successfully. We would like to thank the following people for their support in developing and reviewing the material included.

A special thanks to Michael Haag of the VMware Storage and Availability Business Unit Product Marketing team for pushing this book over the finish line and making it a reality.

Thank you William Rogers, Elinor Schott, Nathan Kawanishi, and Jason Lee for your efforts in completing this book.

Thank you Duncan Epping and Peter Keilty without whose insights this book would never have come to fruition.

And finally, a big thank you to Clinton Lee, Jase McCarty, Jeff Wong, Ken Osborn, and Michael Hall for your input and review leveraging your wealth of experience working with customers who are implementing vSAN.

Preface

Operationalizing VMware vSAN® offers guidance to management-level decision makers and influencers concerned with the impact of VMware vSAN on their current organization, talent, and operational processes as well as engineers and technical decision makers responsible for optimizing operational tooling in a VMware vSAN environment.

Operationalizing VMware vSAN® provides the information needed to help optimize the on-going operations of a VMware vSAN environment considering the critical aspects of team structure and culture, roles, responsibilities, and skillsets, and operational processes as well as monitoring and troubleshooting in a VMware vSAN environment.

Foreword

In my day to day role I spend the majority of time talking about features that VMware vSAN provides today and will be offering in the upcoming years. This is what most architects and administrators prefer to talk about; technology. Being involved with vSAN since the early days I have come to realize however that the introduction of hyper-converged infrastructure software is very similar to the introduction of a virtualized platform, or Software Defined Data Center for that matter. People, and organizations, go through the same phases. They get acquainted with the solution by using it for a very specific use case, this can be for instance virtual desktops or maybe a management cluster, then expand to other parts of the data center.

The key design goals for vSAN were reliability, consistency and simplicity. The technology as such is relatively easy to design and implement, but that is only a small portion of it. The majority of time is typically spent operating the solution. Hyper-converged infrastructure is no different in that regard, and vSAN is the key enabler for HCI. Even though vSAN is relatively easy to manage and monitor compared to a traditional storage system, it is significantly different. Not only from a technology stance, but even more so from a people and process point of view.

As functionality moves to a different layer in the stack, responsibilities will need to be reviewed. As layers of your traditional infrastructure converge, roles within your IT organization may converge as well. Change in responsibility, and even roles, is typically the hardest aspect of the transformation from traditional infrastructure to a hyper-converged based software defined data center. At a bare minimum siloes will need to be broken down and processes will need to be reviewed and possibly revised. Very similar to when virtualization was introduced in to your infrastructure.

Operationalizing VMware vSAN is the perfect companion to go through this exercise. Kevin Lees and Paul Wiggett take you by the hand as they walk you through the considerations and steps required to introduce changes in your organization, and operational model, that will enable you to operate as a provider of IT services. Kevin and Paul are both experts in the field of operational readiness. They have been leading some of the largest projects in the world where an operational (people and process) transformation was required to fully harvest all the benefits of a software defined data center.

I hope you will enjoy reading this book as much as I did.

Duncan Epping
Chief Technologist
VMware - Storage & Availability

Abstract

As a key component of VMware's approach for helping customers evolve to hyper-converged infrastructure, VMware vSAN™ provides an agile, secure, and high-performance solution ready for future hardware, cloud, and application changes. But, to fully leverage vSAN's capabilities in a sustained manner, VMware highly recommends making adjustments allowing you to optimize the way you operate a software-defined infrastructure. This book intends to provide you with guidance and best practices for achieving that optimization when operating a vSAN- based storage infrastructure. It not only addresses tactical optimizations such as monitoring and troubleshooting but also those of a more strategic nature such as the people and team considerations as well as supporting ITSM process considerations. By adopting vSAN technology as well as adapting your operational model, you will fully realize the benefits provided by, and leverage your investment in, software-defined storage.

Introduction

The Software-Defined Data Center (SDDC) presents an almost unprecedented opportunity for IT to dramatically increase its agility and improve time to value in support of business initiatives. In the first book of this series, we covered operationalizing VMware NSX®, which provides software-defined networking and security. This book covers how to operationalize VMware vSAN™, VMware's vSphere-native storage solution which powers industry-leading Hyper-Converged Infrastructure (HCI) solutions with a resilient, high-performance architecture. While both vSAN and NSX are key building blocks for the Software-Defined Data Center, they are merely enabling technologies. As with NSX, simply implementing vSAN does not guarantee your IT will suddenly become more agile or increase the speed with which it provides the business what it needs.

What do we mean by "Operationalizing vSAN"?

"Operationalizing vSAN" refers to what happens after you've designed and implemented vSAN in your infrastructure. You may think, well isn't "day 2 operations" what happens after we've designed and implemented vSAN? That's true, but to really leverage the capabilities provided by vSAN as well as the Software-Defined Data Center and related software solutions, you need to think beyond just day 2 operations.

Some important questions to ask are: How do you optimize the way you use vSAN? Are you currently organized in such a way as to fully take advantage of what vSAN provides? What skills are needed? How can you improve your current operational processes? How might they impact the benefits you could realize from the software-defined nature of vSAN? How can you realize the full benefits of vSAN? What operational tools can help you get the most out of vSAN? What additional considerations does vSAN bring to your IT environment? Can you leverage existing investments in VMware tools? The answers to these questions are what we mean by Operationalizing vSAN.

Why you should adjust the way you operate for SDDC/vSAN

As a pioneer and leading innovator in "software-defined" computing, VMware is uniquely qualified to help organizations in their digital transformation and enable their digital business. To be successful in their digital transformation, organizations must shift to a service mindset, which means aligning their business priorities and IT imperatives. To enable this, IT must also transform. According to a CIO paper ("How IT Organizations Can Achieve Relevance in the Age of Cloud[1]," 2013), "Instead of cost centers that provide capabilities, IT organizations must become internal service providers supplying business-enabling solutions that drive innovation and deliver value… true business partners rather than increasingly irrelevant, cost-centric technology suppliers." To accomplish this shift, IT must consistently and continuously deliver value to the business at the speed the business requires. VMware's software-defined approach provides the technical capabilities to enable the agility and speed with which IT changes can be made and services delivered. Implementing these enabling technologies provides IT the opportunity to change the way it

1 https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/whitepaper/achieve-relevance-in-age-of-cloud.pdf

works. By changing the way IT works to leverage vSAN and Software-Defined Data Center capabilities, you can unlock IT's ability to quickly respond to changing business needs and provide increased business value.

As a result of working directly with Software-Defined Data Center customers, we found that changing the way IT works should be an evolutionary journey across three dimensions to fully leverage software-defined capabilities: People, Process, and Tooling, as shown in Figure 1.1.

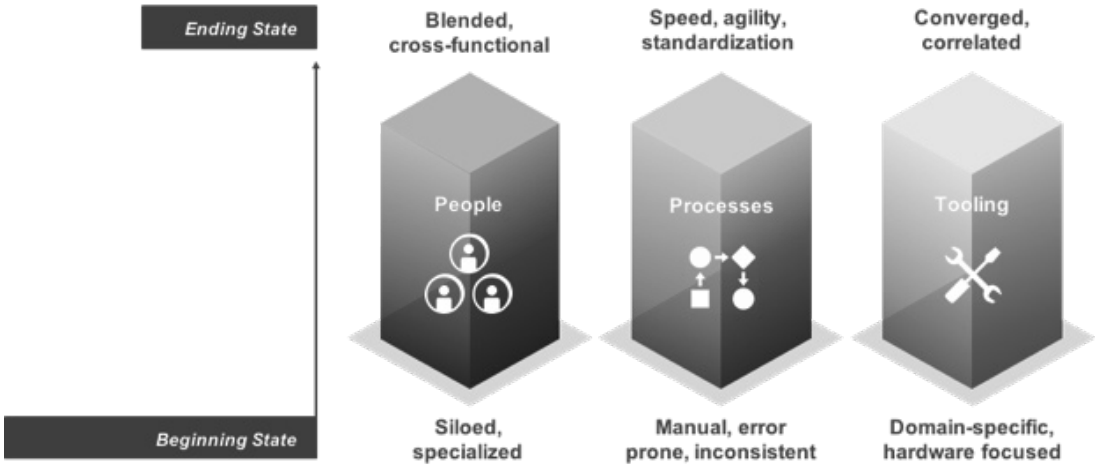

Figure 1.1 Typical beginning state and target end state

Organizations using the traditional IT operating models with siloed functional teams, manual processes, and domain-specific, component focused tooling struggle to successfully operationalize their Software-Defined Data Center, and this applies equally to vSAN. Working across these three dimensions to break down the siloed teams, automate and optimize processes to take advantage of hyper-converged infrastructure software capabilities, and move towards stack-aware tooling allowed these customers to begin moving with the speed and agility the business needs, and to realize the full benefits of their investment in vSAN. This change doesn't happen overnight, nor should it. It is an evolutionary journey from where you are to where you need to be.

Key Takeaway: Transforming your people, process, and tooling to optimize how you operate your software-defined environment should be evolutionary, not revolutionary.

Who the guide is for

This guide is meant both for people who are new to vSAN as well as those who have implemented vSAN but are looking to get even more out of it.

- Organization level decision makers who are concerned with the people perspective: organizational impacts, team structure, roles and skillsets, culture and mindset
- Managers responsible for or who can influence operational processes such as change and configuration management
- Architects and technical decision makers responsible for optimizing tooling to be more effective in a software-defined environment

What you will learn

You will learn what you need to do to operationalize your vSAN environment at an introductory level. The goal is to get you started, and to provide references to more detailed documentation.

Why you should care

Your company has invested in, or is looking to invest in, vSAN and hopefully the broader SDDC suite. The technology enables you to provide greater value to your users, whether IT itself, developers, application owners, or ultimately your company's end-users. It's presumably in your best interest to derive the greatest value you can from that investment. We've found through working with many organizations, to unlock the business value of vSAN and SDDC, you need to act on the guidance provided in this book. Those who have, increase their value not only to their companies but to their careers. The software-defined data center is the future. Think of this book as a way to ensure you remain relevant, and to increase your personal value to your company and beyond.

How to proceed

Different sections are more valuable for readers who are coming from different perspectives. Organization level decision makers would benefit from the entire book for context and increased understanding of how to unlock vSAN's potential. If you're short on time, focus on "Measuring Results" and "People Considerations" first, then "Intelligent Operations." If you are a manager interested in optimizing your operational processes, focus on the process related sections in "Intelligent Operations." If you have a technical bent, focus on the monitoring and troubleshooting sections of "Intelligent Operations," as well as "Tools for Monitoring and Troubleshooting."

Measuring Progress

You've probably guessed by now that we're encouraging you to make some changes that will allow you to fully realize the benefits of vSAN. In other words, don't try to fit vSAN into your current way of doing things, that's not a path to optimal value realization. As with any change, especially transformative change, you need to know that you're progressing and what results you're attaining.

Why measure

Let's get right to the point. When you undertake the effort to change, you need to measure that change so you can prove to the decision makers you're making demonstrable progress and achieving a return on their investment. It's more powerful if you can back up your talk of success with data. Telling your executives you've been able to lower your total cost of storage ownership by 50% with capital and operational savings, as well as see a 10% increase in productivity due to workload-attached storage policies, will have a lot more impact than saying "yep, this new hyper-converged storage is great!" Ok, this may be a bit extreme (though we've seen it done to great effect), but you get the idea.

In addition to impressing your executives, it's also good to measure progress and results to drive conversations that help teams improve. Notice here it's about improving, not evaluating team performance; and providing positive reinforcement to encourage continuous improvement.

How and what to measure

You need to start by having a good understanding of your current state of operational performance. Establish a baseline against which to measure improvement and ultimately claim success in your efforts. We provide some key performance indicators in each section of the book. Take these key performance indicators as a starting point and measure how you're currently performing against each.

Finding quick wins and with whom to share results

Quick wins are key. Don't try to "boil the ocean" by making a bunch of changes at once. Start by focusing on processes that will deliver the most value to IT or your end users, but require the least amount of effort. Choose a specific application or service to focus on initially and create a tiger team following the guidance provided in "People Considerations" section. Perhaps you could begin by creating VM storage policies and applying them to tiers of a specific application. How long did it take you to perform those same storage assignment steps in the physical world? How long would it take to purchase, burn-in, configure, and deploy storage for a new application? Compare that to how long it took to define and deploy the required VM storage policies for vSAN using Storage Policy-Based Management (SPBM). What were the savings in time and money from both the perspective of pure deployment time, as well as cost savings from both a people and

licensing perspective? The answers to these questions will help make your progress goals more real.

It's good practice to share your results up your management chain, down your management chain, and laterally to other teams within IT! This is a new technology and a new way of working as it relates to storage. You want to reinforce up the IT management chain how your hyper-converged storage implementation is producing valuable IT and business outcomes. You want to further incentivize success through acknowledgement and recognition down your management chain. You want to market your success laterally across IT to both quiet the naysayers, as well as inspire others to get involved as application of vSAN expands. Finally, don't forget about your business stakeholders; be sure to actively communicate and market your success to them – if it directly resulted in providing a demonstrable business outcome and you can put it in business terms to which they can relate.

Key Takeaway: Focus on quick wins and share the results up your management chain, down your management chain, and laterally to other teams within IT.

People Considerations

Whether you're a smaller IT organization just getting started or a larger IT organization focusing on truly becoming a service provider to your business stakeholders, you're focused on ultimately delivering business value. vSAN, like the Software-Defined Data Center, is an enabling technology, but to fully realize its potential for contributing to the business value you are striving to provide, we recommend that you optimize your operating model. The starting point is the people perspective. Aligning team structures as well as roles and skillsets, and affecting cultural and mindset changes are the basis for these changes.

Why SDDC (and vSAN) changes the people equation

Operationally, IT organizations are traditionally composed of technically-aligned functional groups. To date, they have been able to get by even if they have a significant investment in compute and memory virtualization-based infrastructure environments. These organizations have certainly gained efficiencies and seen improvements in virtual workload deployment times by automating components of their virtual workload deployment process, but those gains are incremental at best. Regardless of the activities in a virtualized environment, if multiple functional teams are needed to complete a task, such as deploying a new application or service, such efforts cannot quickly respond to changes in the business's requirements or quickly deliver what the business wants.

The challenge doesn't stop with technically-aligned functional teams. IT organizations of any size are inevitably organized functionally in plan, build, and run siloes. Architects make technical decisions and design solutions; engineering teams build and test the solutions architects provide them; and the solution is handed to IT operations to run. How long does this end-to-end process take? More importantly, how prepared is your IT operations to run it?

This can even extend beyond putting full solutions in production, which can be even more problematic. One customer applied a subset of this process to developing and modifying VMware vRealize® Operations Manager™ dashboards. Operations had to provide requirements to a tools engineering team who, due to the inevitable backlog, provided the dashboard some two months later. Of course, once the operations team received the dashboard, assuming they remembered they even requested it, it didn't provide them what they needed!

By providing in software what was once only hardware-based, vSAN and the Software-Defined Data Center enables fully virtualized compute, memory, network, security, and storage infrastructure solutions, and the workloads running on them to be implemented and changed very quickly. Add the automation opportunities and IT has the potential to provide previously unheard-of levels of business value. This potential will never be reached using existing IT functional grouping constructs. Maintaining these constructs will also prevent IT from ever adopting the Agile methodologies being effectively leveraged in application development; methodologies that can lend themselves to IT being even more responsive to and faster in delivering solutions that meet changing business needs.

Key Takeaway: You have to breakdown siloes across two dimensions to be successful: plan-build-run and technically-aligned functional teams.

Team construction

vSAN changes storage management and consumption dynamics. It frees up storage administrator's time to focus on the backend tasks they're so good at, like storage capacity management and expansion, performance tuning and troubleshooting, and enhancing the quality of service. vSAN does this by relieving them from being continually interrupted by mundane tasks such as LUN creation and reclamation or file system management. However, to fully leverage vSAN capabilities in the context of a Software-Defined Data Center comprised of a fully virtualized stack, you might want to optimize your team structure. While these optimizations aren't critical or even necessary when considering vSAN as just another component in your overall storage environment, they become advantageous when you consider the bigger picture. It's equally true when considering vSAN in the context of VMware Cloud on AWS or Hyper-Converged Infrastructure powered by vSAN. The goal of these optimizations is a blended or integrated team; a team consisting not only of cross-functional technical skills, but also cross-domain roles aimed at creating a much closer relationship between architecture, engineering, and operations for a Software-Defined Data Center-based environment that includes vSAN. This blended team serves as the focal point for all decisions and actions regarding the environment.

The goal in creating a blended team is to break down the cross-functional and cross-domain IT silos, shown in Figure 3.1, that inhibit agility and execution speed. It intends to replace these silos with a team built for tight collaboration and focus. Creating such a team results not only in faster but better decision making, reduced time to resolution of problems, and more operationally ready solutions.

Figure 3.1 Cross-domain and cross-functional siloes

How does a blended, cross-domain team bring success with vSAN and SDDC? In our experience working with organizations, you can no longer afford to have distinct plan, build, and run teams if you want to maximize the agility afforded by software-defined infrastructure. Operations needs to have more involvement in architecture and design decisions. A common example involving vSAN is: The Plan team chooses vSAN (great choice!), the Build team implements it (so far so good), but then it's "thrown over the wall" to the Run team who was not part of the plan and build process. As a result, when they start running vSAN they're rebooting servers and taking other actions that are appropriate generally in their virtualized environment, but ones that can wreak havoc in a vSAN storage cluster. A tight feedback loop is needed between plan, build, and run when working with a software-defined infrastructure. We have seen that organizations who adopt these cross domain functions as part of a blended team are able to more quickly and efficiently scale their software defined environments than those who did not. A typical blended team is shown in Figure 3.2 below.

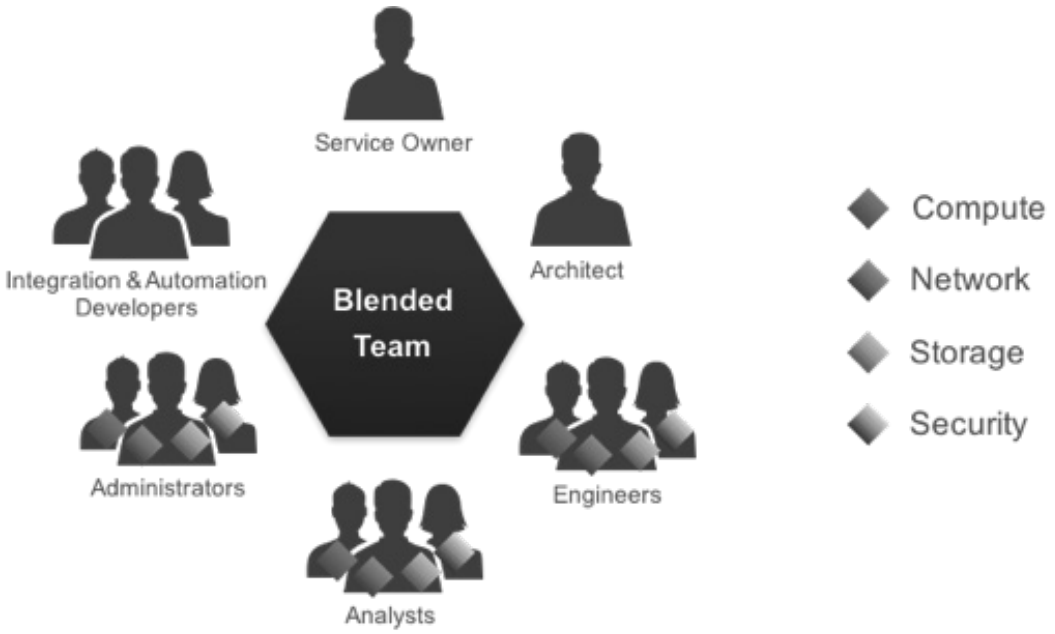

Figure 3.2 Blended Team

The same is true for breaking down the technical silos and creating a blended, cross-functional team, which is also shown in Figure 3.2. Relying on silos of virtualization, network, security, and storage to accomplish a result in a software-defined infrastructure is antithetical to achieving agility and speed of execution. A great example comes from vSAN's dependence on connectivity between nodes to provide access to the distributed storage resources. The fact that connectivity can be permanently or intermittently disrupted is bad enough, but

changes made by the network team completely unbeknownst to a virtualization team or storage team simply because they're siloed should not be one of them. We're seeing exactly these types of issues arising all too often when operating software-defined environments; issues that are addressed directly through a blended team approach. Siloed teams must be left behind to be successful in a software-defined data center.

This blended team can be physical or virtual; matrixed or under a single manager. There is no best answer as we've found these decisions tend to be very company-specific. They often involve politics and cultural dynamics. Achieving a blended, cross functional, cross-domain team is what matters in the end. We've seen it work in various organizational and reporting structures. The common success factors are: the team is built with a clear purpose and shared objectives; they're incentivized on achieving team objectives more so than individual objectives; and they're as self-sufficient as possible with decision authority, and with the associated accountability, to achieve their objectives. In addition to technical and domain skills, what's important is to have the team consist of members who are change agents, enthusiastic, and can naturally act as evangelists for what they're doing. This is a new technology and way of working. Team members must embrace this change and make it infectious to others. This is critical to ensure on-going success as hyper-converged storage as well as the software-defined data center, starts to represent a bigger percentage of the overall infrastructure and becomes more business critical.

Key Takeaway: Forming a blended, cross-functional, cross-domain team is paramount; whether it is virtual, physical, matrixed, or under a single manager is a company-specific organizational decision.

Creating such a team and determining its size is circumstance dependent. Are you just getting started with a new vSAN implementation? Are you adding vSAN to an existing virtualized environment or is this part of a larger, "greenfield" software-defined data center implementation? Are you experimenting with hyper-converged storage or is this a situation where the company has made a strategic decision to go with a software-defined data center including vSAN?

If you are just getting started with a new vSAN implementation or "dipping your toe" into hyper-converged storage or perhaps there just isn't yet the management buy-in to create a permanent blended team, we've found following a process like that shown in Figure 3.3 to work well.

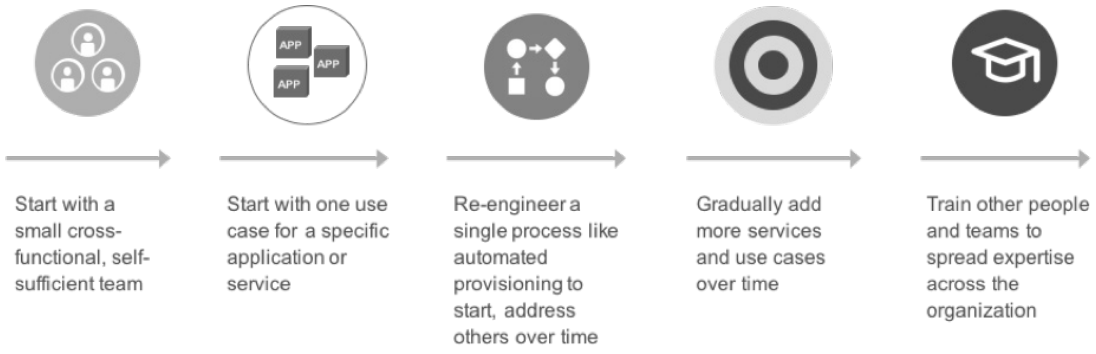

Figure 3.3 Getting Started

We've found that taking a "tiger team" approach works best. For example, create an initiative focused around some aspect of IT automation for example. A common example is automatically applying storage policies to three-tier applications provisioned using VMware vRealize® Automation™. To do this, you want to define and create a blueprint that can be standardized and used repeatedly. This type of effort requires input from the application team and all the virtual infrastructure teams: compute, storage, network, and security. And to truly control the resulting automatic provisioning service end to end – plan, build, and run – it requires involvement from architecture, engineering, and operations. Pull together a cross-functional, cross-domain tiger team, even if it's temporary, to drive your consensus-building initiative. Remember to make part of the initiative to baseline current and resulting metrics around provisioning storage for these three-tier application infrastructures – both to prove to yourself the benefits of a blended team as well as to put forth a data-driven argument for making it the norm going forward. Keep the decision-makers informed from the beginning of the initiative, provide them with regular updates, and schedule a final read-out session in which you can demonstrate the value of a blended team based on the data-driven results you realize.

Key Takeaway: If you are new to vSAN or the software-defined data center, start with a small, blended tiger team focused on a single use case for a specific application or service and expand over time.

From a functional perspective, compute, network, storage and security should all be represented on the team responsible for the software-defined data center. The initial tendency we see is for organizations to simply have the team responsible for virtual environments take on all software-defined components – after all it worked for compute didn't it? While it might be easier to see the challenge in doing so for network and security, it may be less obvious for storage. For example, expecting a virtualization administrator to assume responsibility for vSAN without understanding storage policy implications can prove

disastrous. If storage policies are changed without understanding the impact on capacity or even compute resources, it can quickly bring down the entire vSAN cluster. This example of a correctable, but easily avoidable, situation shows the need for a blended team, one in which cross-functional skills are needed for success.

Creating a blended team is critical to fully leveraging vSAN capabilities and achieving game-changing success in the context of a software-defined data center. Everything else drives from this blended team structure. If you cannot get buy-in to create a blended team, you will need to manage expectations of dramatic improvements in agility or speed of execution in meeting your business stakeholders' changing needs. In this case, our advice is still to go through a tiger team exercise while measuring results against a baseline. Use this to champion a strategy of pulling together temporary tiger teams to drive business impacting improvements using your hyper-converged storage capabilities. Following this approach can provide incremental improvements, but the challenge will be sustaining the improvement operationally. The hope and goal must be gaining mindshare for creating a blended team, ultimately taking this approach.

Roles and skillsets

This section provides guidance on the recommended roles for storage in a hyper-converged environment. It also describes key responsibilities, skillsets, and education associated with each role. We focus on the storage aspects of hyper-converged infrastructure, not all the roles we would recommend for a fully blended team. For more information regarding the other roles in a blended team please download the "Organizing for the Cloud" whitepaper referenced in Table 8.1 in the "Where you can go for more information" section.

Keep in mind that we're discussing roles in this section, not headcount. In small and mid-size environments, roles may be combined into a smaller number of individuals whereas in large, business critical environments we've seen multiple individuals providing a single role perhaps with even more specialization.

Our approach to a fully blended team in the broader context of the software-defined data center tends towards full stack knowledge with some specialization. While we're addressing the roles only in the context of hyper-converged storage, we do provide a description of how the role is impacted in our approach to fully blended teams for the software-defined data center generally.

While we describe a Cloud Storage Architect role, we favor an overall Cloud Architect role that may have more general storage, as well as network and security skillsets for a software-defined data center blended team. In that case, we focus the storage-specific skills in the Cloud Storage Engineer role. He or she works closely with the overall Cloud Architect, providing the needed storage subject matter expertise.

Table 3.1 Architecture Roles

Role	Responsibility	Skillsets and Education
Cloud Storage Architect	• Identify and prioritize use cases and business requirements to address with hyper-converged storage • Design logical storage services for availability, capacity, recoverability, and data protection • Design standards and templates for automated hyper-converged storage infrastructure provisioning • Verify the hyper-converged storage infrastructure by developing and validating tests to ensure the success of addressing use cases and requirements • Identify modern tools for hyper-converged storage infrastructure automation, and day-2 operations (visibility, monitoring, troubleshooting) • Guide the hyper-converged storage infrastructure implementation strategy • Provide level 3 support as needed to work within defined SLA or OLA resolution period • Assist in defining and evolving overall IT storage architecture and standards that maximize the synergy, reuse and value of vSAN over time • Proactively engage with VMware field resources to understand vSAN's product direction and roadmap	• Cross-domain skills (i.e., hyper-converged storage, network & security, vSphere, vSAN availability and performance characteristics) • Education: Data Center Virtualization Fundamentals; vSAN deploy & Manage; vSAN: Troubleshooting Workshop; vSAN Production Operations; vSAN Hands-on Labs • Certification: VCDX-DCV or VCDX-CMA

The Cloud Storage Engineer role is key, whether this is a vSAN-focused blended team or a broader software-defined data center-focused blended team. Whether providing deeper storage subject matter expertise to the architect roles or designing storage policies implemented by the administrator roles, the storage specialization is focused in the engineer role.

We've also included the Cloud Automation & Integration Developer role in the table below, as it is such a critical role in the modern, software-defined data center world. This role is critically important in the context of a vSAN-focused blended team or a broader software-defined data center-focused blended team. Automation is absolutely key to success going forward and automation-related workflow development must move beyond administrators writing scripts to something more akin to actual software development. This is where the Cloud Automation & Integration Developer comes into play.

Table 3.2 Engineering Roles

Role	Responsibility	Skillsets and Education
Cloud Storage Engineer	• Low-level design, deployment, and testing of the hyper-converged storage functions that realize the hyper-converged storage infrastructure, definition of the hyper-converged storage infrastructure function configurations, ensure that the hyper-converged storage infrastructure services and functions work successfully, and make them operational • Supports the Cloud Storage Architect role in designing cloud storage services and translates the requirements into blueprints and storage policies for the storage functions realizing the storage service • Ensure fulfilment of requirements, including capacity, availability, performance, security, compliance, SLAs • Deploy, test, validate, and manage monitoring and troubleshooting tools, processes, dashboards, runbooks • Works with the Cloud Automation & Integration Developer role to design, develop, test and deploy custom workflows and scripts within the hyper-converged storage infrastructure that can be used for integration, orchestration, deployment, monitoring, compliance, or other routine tasks. • Provide level 3 support as needed to work within defined SLA or OLA resolution period • Diagnose and analyze root cause of issues, and work with Cloud Storage Administrator to apply patches and fixes as needed • Implement routine, approved and exception changes in the hyper-converged storage infrastructure • Assess and test upgrades and patches for hyper-converged storage infrastructure and tools • Proactively engage with VMware field resources to understand vSAN's product direction and roadmap as well as online resources and the local technical community to stay abreast of vSAN recommended practices	• Cross-domain skills (i.e., hyper-converged storage, network & security, vSphere, vSAN availability and performance characteristics) • Education: Data Center Virtualization Fundamentals; vSAN deploy & Manage; vSAN: Troubleshooting Workshop; vSAN Production Operations; vSAN Hands-on Labs • Certification: VCDX-DCV or VCDX-CMA

Role	Responsibility	Skillsets and Education
Cloud Automation & Integration Developer	• Works with the Cloud Storage Engineer to establish integration and automation monitoring • Works with the Cloud Storage Engineer to establish automated storage service provisioning as well as event and incident remediation wherever possible and appropriate	• Skills: PowerShell, PowerCLI, Python, Ansible), Configuration Management tools (Chef, Puppet), Orchestration tools (VMware vRealize™ Orchestrator™), vSAN API, vRealize Automation API, VMware vRealize™ Operations™ API • Education: Data Center Virtualization Fundamentals, Data Center Automation with vRealize Orchestrator and PowerCLI, VMware Cloud Orchestration and Extensibility, vRealize Orchestrator: Develop Workflows

Finally, we have the Cloud Storage Administrator role. In addition to Day Two operations like backup and restore, upgrade and patching, the administrator role is focused heavily on proactively monitoring and remediating the hyper-converged storage infrastructure. He or she is also responsible for working with the engineer and developer roles to customize the monitoring tools to continuously improve their proactive and predictive capabilities. The goal is to minimize the actual number of incident tickets received by proactively identifying and remediating issues before they become service or application disrupting. While this is typically fulfilled by the same person administering vSphere, it's very important he or she has the appropriate level of vSAN-specific understanding and skills.

Table 3.3 Administrator Roles

Role	Responsibility	Skillsets and Education
Cloud Storage Administrator	• Monitor hyper-converged storage infrastructure and act on events before they affect services • Proactively monitor vSAN performance (latency, throughput), health (faults, failures, connectivity), availability, and configurations • Update and maintain hyper-converged storage infrastructure by utilizing alarm/alert mechanisms • Provide level 3 support for hyper-converged storage • Investigate and diagnose hyper-converged storage infrastructure and services incidents • Ensure solutions and fixes are applied to recover from storage incidents • Implement and apply storage policies designed and tested by the Cloud Storage Engineer • Backup and restore of vSAN manager data (e.g., system configuration, events, audit log tables) • Upgrade and patch hyper-converged storage infrastructure and tools • Define, develop, and test custom dashboards, super metrics, reports, etc., for hyper-converged storage infrastructure and services; work with the Cloud Automation & Integration Developer role to implement remediation automation capabilities • Proactively engage with VMware field resources, online resources and the local technical community to stay abreast of vSAN recommended practices.	• Cross-domain skills (i.e., hyper-converged storage infrastructure, vSphere, storage policy considerations) • Education: Data Center Virtualization Fundamentals; vSAN Deploy & Manage; vSAN: Production Operations; vSAN: Troubleshooting Workshop; vRealize Operations for Operators; vSAN, vRealize Operations, & VMware vRealize® Log Insight™ Hands-on Labs • Certification: VCP-DCV or VCP-CMA; VMware vSAN Specialist (Badge)

Just to reiterate, we described roles in this section, not headcount or individuals. Oftentimes, depending on the scale and business criticality of the environment, and the level of automation, multiple roles may be filled by a single individual, or multiple individuals may fill a single role. The most important point is that these roles and skillsets exist in the blended team.

Key Takeaway: Focus on full-stack knowledge with some specialization when filling roles and putting together training plans for your blended team.

Culture and mindset

Culture and mindset can be thought of as the equivalent of the organization's DNA – the values and beliefs that shape how people behave and create the organization's culture. In our experience, this is the single most impactful factor influencing success or failure for your adoption of hyper-converged storage and Software-Defined Data Center. It's also the most difficult to change. This is one of the first things we assess when organizations with which we work get serious about implementing vSAN.

Ideally, the culture we want to see is one that embodies collaboration, a business outcome focus, shared objectives, system thinking, continuous learning and continuous improvement, driven by teams that are both responsible and accountable for achieving objectives they acknowledge as owning. These are many of the characteristics associated with the Agile culture popular in application development teams. This represents a rare breed of organization, and one that we rarely see. Achieving this type of culture, though, is a critical goal in order to be successful.

If this doesn't describe your immediate organization, how can you move towards it? There are entire books devoted to answering this question, but some guidance we can provide based on our experience includes:

- Leadership: Cultural change must be embraced from the top down even if it is only realized in the vSAN-related, or software-defined data center related blended team
- Direction: Leaders must articulate a clearly defined direction for the blended team and modify incentives to reinforce shared team objectives
- Ownership: The team must be given ownership, responsibility, and accountability for achieving the stated purpose and shared objectives of the team
- Change Agents: The members initially selected for the blended team must be like-minded change agents; they must be open-minded and passionate about instituting the culture within the team
- Recognition: Team successes must be recognized and advertised; the goal is to make their behavior aspirational for others in IT. What the team is doing and how they're progressing should be actively marketed within IT and, as appropriate, the key business stakeholders of the software-defined environment

Key Takeaway: **Overcoming cultural challenges is always the biggest hurdle and should be explicitly addressed from the beginning.**

Sample team structure KPIs

Key Performance Indicators (KPIs) are used to measure progress towards target state objectives which, in turn, should clearly communicate the measures of success. Of course, target state objectives are company or organization specific and depend on their goals, ideally tied to IT and business outcomes, for implementing hyper-converged storage. The following are some KPIs to begin measuring progress towards establishing a blended team generally for a software-defined data center. Also note, we prefer KPIs based on comparison against a baseline – plus this has the added benefit of reinforcing the need for creating a baseline against which to measure.

Table 3.4 Sample team structure KPIs

Objective	KPI	Description
Self-sufficient team	• % of escalations resolved within the blended team • Average time to resolve escalations with the blended team	These KPIs provide a measure of blended team efficiency in resolving issues versus the baseline of how long it previously took to resolve escalations across siloed teams.
Shared objectives	• Team-based annual review criteria as a % of team member's review criteria • Average time to complete cross-functional activities within the blended team	These KPIs provide a measure of a blended team's efficiency in completing cross functional activities (such as on-boarding a new application) when they have team-based objectives versus the baseline of completing a similar activity involving siloed teams with competing objectives.

Process Considerations

Intelligent operations

Most IT organizations can't seem to break out of "firefighting" mode and are constantly monitoring and reacting to events and alerts being generated in their environment. They're also plagued with laborious operational tasks consisting of error-prone manual steps. Applying this mode of operation to a vSAN-based environment will never allow you to realize the full benefit of hyper-converged storage and the software-defined data center. The reality is, working in this way it is most likely that you will never be able to realize the agility and speed of execution offered by the SDDC and its enabling technologies such as vSAN and NSX.

Intelligent Operations refers to a proactive mode of operation rather than the reactive mode of operation that results in constant "break-fix" activities, and the feeling of never having the time to innovate and make educated improvements in your environment. Intelligent Operations is also about optimizing and automating your processes and work habits to take advantage of the software-defined data center. It's about using the right tools for the job; tools built for Intelligent Operations in a software-defined infrastructure. This can span from the automated lifecycle management capabilities of VMware's hybrid cloud platform, VMware Cloud Foundation™, to taking full advantage of policy-based provisioning, placement, and operational capabilities in VMware vRealize® Suite, as well as leveraging intelligent analytics engines in tools like VMware's vRealize Operations Manager.

This chapter aims to outline the process considerations we need to take into account when performing Intelligent Operations as it relates to hyper-converged storage. In this chapter, we'll focus on those processes and activities most impacted by vSAN and how they can be optimized and, in some cases, automated. Of course, we work for VMware so the emphasis and examples will be based on VMware's native tools. We will be sure to describe the suggested tools in terms of capabilities, so you can substitute these with other tools, providing the same or similar capabilities. VMware has worked extensively to establish deep integrations and partnerships with tooling platforms with which you maybe already invested, as well as an extensive API set for custom integration.

Proactive monitoring for performance and availability

Proactive monitoring is more a cultural or mindset change than about technology. As a key to success in becoming a service provider, this is about shifting your approach to proactively monitoring key metrics, analyzing potential issues, and remediating those issues before they become service, application, or end-user impacting. Even though we're focused on hyper-converged storage, proactive performance and availability monitoring is about putting more emphasis on monitoring from a service- or application-centric perspective. Where, for example, a service may be Database as a Service provided to application developers or Data Analytics as a Service being consumed by lines of business. While proactive monitoring depends on a mindset shift to a different approach, the selection of monitoring tools matters in that it not only enables the ability to perform proactive monitoring, but also dictates how effective you can be in achieving it successfully.

What does it mean to put more emphasis on monitoring from a service or application-centric perspective? First and foremost, this means monitoring the full stack on which the application or service is running, and doing it in an integrated way. This means monitoring the virtual machine or set of virtual machines making up the application or service, the hypervisor and host, storage, and the end-to-end network path underlying the application or service; in other words, anything that impacts application or service performance or availability. Including integrated network monitoring is particularly important, given vSAN's network dependencies. And in the case of performance, it includes monitoring the usual suspects like CPU and memory contention along with, in the case of vSAN, the aggregate of read and write latency, when considering any performance related service level agreements associated with the application or service.

If a service level agreement is involved, it's best practice to set your monitoring thresholds at some percentage below the service level threshold, to provide a buffer. This is a great example of proactive monitoring; by setting a threshold at some percentage below the service level threshold, you have bought yourself some time to troubleshoot, if performance or availability is consistently hitting the lower threshold, before something becomes application or service impacting. Using vSAN performance as an example, if your aggregate latency is 10 ms for your highest service tier, you might start by giving yourself a 25% buffer by setting your monitoring threshold at 8 ms and adjust accordingly as you collect actual data.

Speaking of performance and availability, it's worth adding a quick check of your top workloads experiencing CPU contention, memory contention, and disk I/O latency to your daily routine – which is easily done with vRealize Operations Manager. vRealize Operations Manager provides both environment and detailed vSAN-specific dashboards to allow quick identification of hotspots and potential issues which need to be resolved. Examples of these are shown in Figure 4.1 and Figure 4.2.

Figure 4.1 Example of vRealize Operations Manager's Operations Overview dashboard

Figure 4.2 Example of vRealize Operations Manager's vSAN Operations Overview dashboard

Table 4.1 Sample proactive performance and availability KPIs

Objective	KPI	Description
Proactive performance monitoring	• Number of performance-related issues detected and resolved before they become service or application impacting incidents • % reduction in performance-related incidents month over month	These KPIs provide a measure of how effective you become at proactive performance monitoring over time.
Proactive availability monitoring	• Number of availability-related issues detected and resolved before they become service or application impacting incidents • % reduction in availability -related incidents month over month	These KPIs provide a measure of how effective you become at proactive availability monitoring over time.

Proactive capacity monitoring and planning

Similar to performance and availability monitoring, proactive capacity monitoring and planning is about shifting your mindset to proactively identifying and remediating capacity issues before they become service, application, or end-user impacting. Unlike proactive performance and availability monitoring, which is best approached from a service and application-centric perspective, proactive capacity monitoring and planning for vSAN focuses mainly on the vSAN cluster. In a vSAN cluster, multiple LUNs and datastores are replaced by a single datastore comprised of the storage contributed by each host in the cluster. This approach also provides flexibility in adding capacity, since you can easily add storage capacity achieved by either scaling up through adding more storage devices or higher density storage devices to each host in the vSAN cluster, or by scaling out through adding hosts to the vSAN cluster. This allows you to build out capacity incrementally and grow capacity in a flexible manner as your business and consumer demand grow. It also means that provisioning of additional capacity is much quicker as lead times on individual disks are relatively short.

Providing sufficient free space in a vSAN datastore is fundamental to maintaining the health of a vSAN cluster (VMware recommends a minimum of 25% to 30% free space). Proactively monitoring available and used storage capacity along with % remaining is a simple but effective way to provide proactive capacity monitoring. vSAN provides out of the box capacity views integrated into VMware vCenter Server to show current high-level status as shown in Figure 4.3. This overview provides a view of total used and free space, as well as capacity savings if compression and deduplication are enabled. It also provides a breakdown of the capacity used as grouped by vSAN object types or data types. Understanding and extrapolating historical consumption rates is key to forecasting when additional storage may be needed.

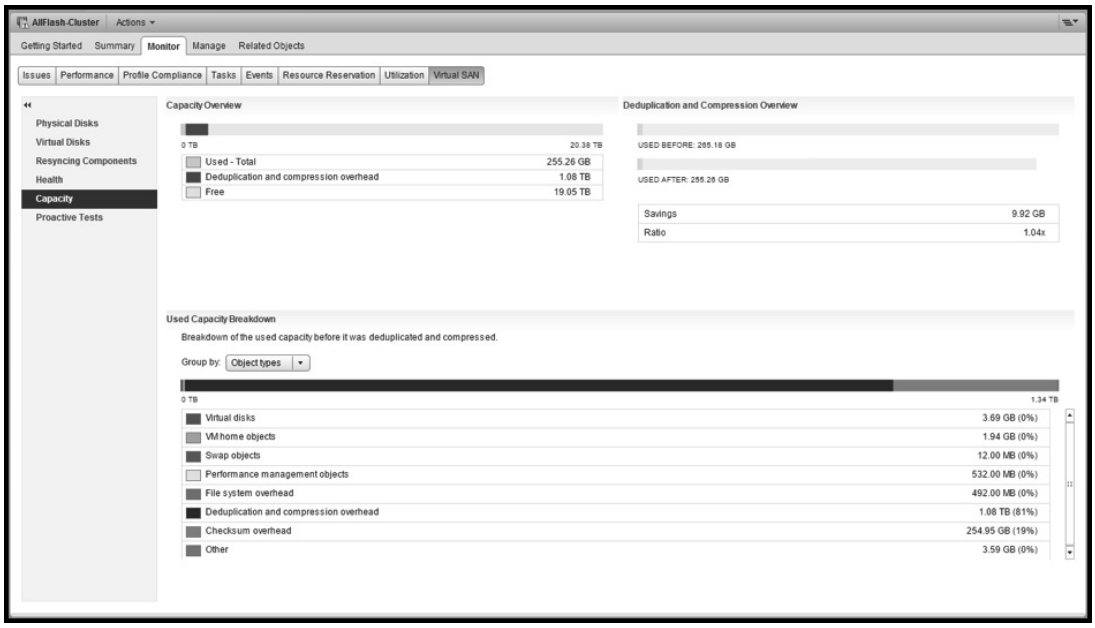

Figure 4.3 vSAN Capacity Overview

Fortunately, tools like vRealize Operations Manager can provide easy to consume, at-a-glance statistics to provide this level of proactive monitoring and planning as shown in the vSAN Capacity Overview dashboard example in Figure 4.4.

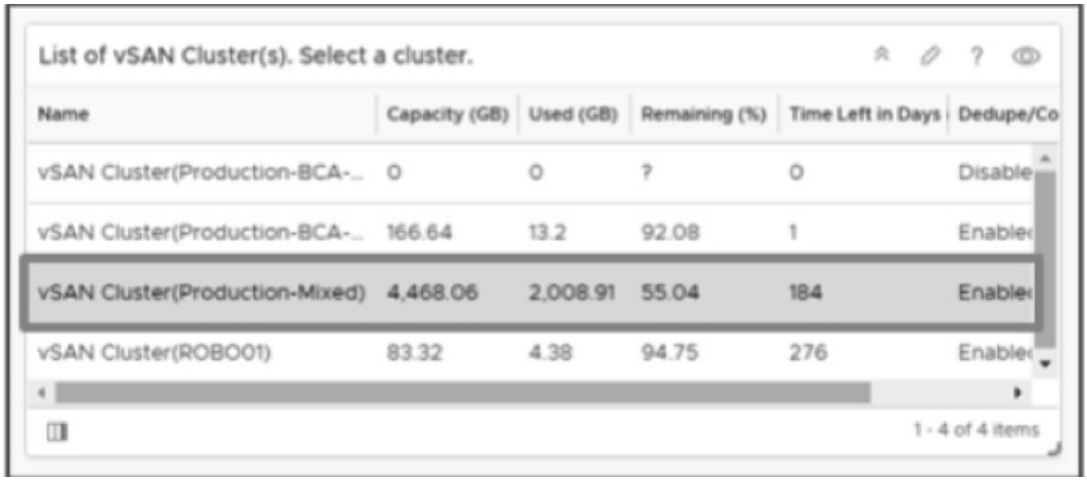

Figure 4.4 vSAN Capacity Overview dashboard

The “Time Left in Days” column extrapolates past consumption rates to forecast when additional storage may be needed so you can plan around the lead time on procuring this as part of your organization’s procurement process. It’s worth noting this is an estimate only and does not reflect operational changes such as deduplication and compression, failure tolerance methods, or the mixture of storage policies impacting storage utilization rates associated with the stored virtual machines on the datastore.

You also should get into the habit of reclaiming unused storage where applicable. Examples of this include storage capacity being consumed by old snapshots and unused virtual machines such as those that are either powered off or idle. This reclaimed storage capacity could realize significant cost savings and efficiency gains especially in larger environments. It also allows you to proactively reduce the number of vSAN components. vSAN currently has a limit of the number of components it can store in a single cluster so proactively managing and reducing this component count is a good practice to pursue, as it also forms part of the proactive capacity management considerations. It is recommended that regular checks be performed every 3-6 months to make sure these reclamation opportunities are identified and actioned. Remember also that any reclamation actions should be validated with the virtual machine owners to make sure the identified object is no longer required. It could be that even though it is idle, it is being used for a disaster recovery function. Figure 4.5 shows an example of virtual machines where disk space can be reclaimed as reported by vRealize Operations Manager. A tool such as VMware vRealize® Business for Cloud™ could then also be used to equate these capacity savings into financial metrics as shown in Figure 4.6. Demonstrable cost savings to business lines are an excellent way to influence positive consumer behavior when attempting to reclaim unused or wasted resources.

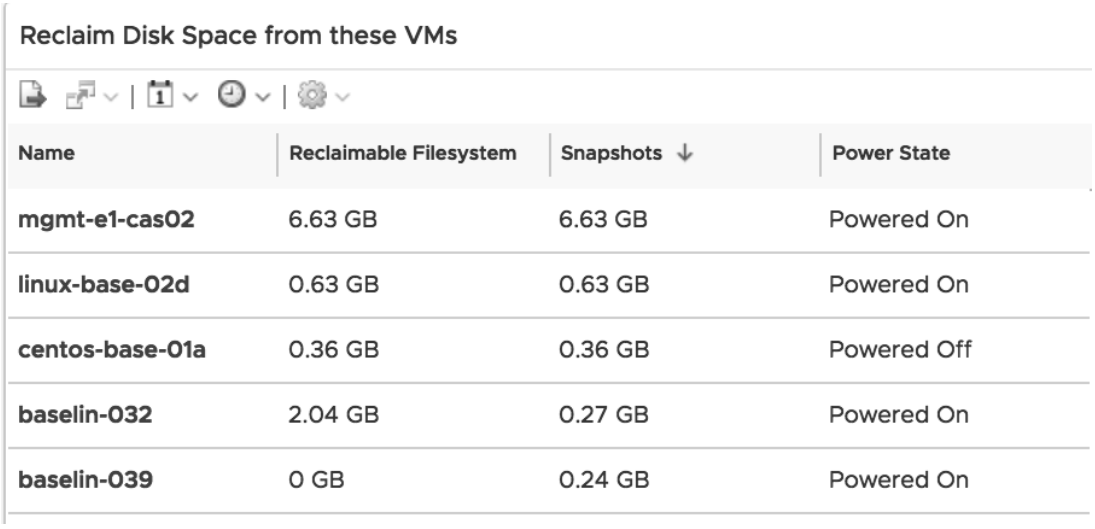
Reclaim Disk Space from these VMs

Name	Reclaimable Filesystem	Snapshots	Power State
mgmt-e1-cas02	6.63 GB	6.63 GB	Powered On
linux-base-02d	0.63 GB	0.63 GB	Powered On
centos-base-01a	0.36 GB	0.36 GB	Powered Off
baselin-032	2.04 GB	0.27 GB	Powered On
baselin-039	0 GB	0.24 GB	Powered On

Figure 4.5 vRealize Operations Manager Capacity Reclaimable dashboard

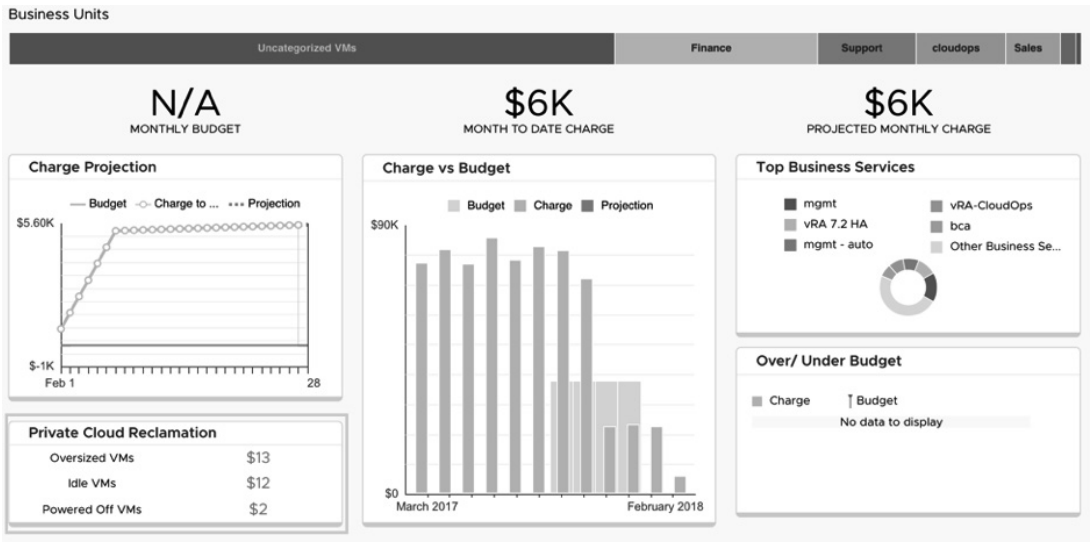

Figure 4.6 vRealize Business for Cloud Showback Statement

Table 4.2 Sample proactive capacity monitoring, planning and reclamation KPIs

Objective	KPI	Description
Proactive capacity monitoring	• Number of capacity-related issues detected and resolved before they become service or application impacting incidents • % reduction in reported capacity-related incidents month over month	These KPIs provide a measure of proactive capacity monitoring effectiveness for any defined vSAN clusters. (Requires Proactive Issue Resolution checkbox along with Capacity Incident Category in ITSM tool)
Proactive capacity planning	• Average amount of time in advance of on-boarding a new application or service that requirements for additional vSAN cluster capacity are identified • Average amount of unused capacity in the vSAN cluster month over month	These KPIs provide a measure of how effective you become at proactively planning capacity needs for new applications moving into the environment as well as how accurate you are in your proactive capacity planning.
Proactive capacity reclamation	• Average amount of idle/unused capacity reclaimed month over month • Cost avoidance from reclaimed capacity	These KPIs provide a measure of how effective you become at reclaiming idle or unused capacity and the potential delay of future capacity investment.

Change management

Highly agile and dynamic software-defined data center-based environments have many different moving parts, and usually contain a large amount of integration points into a variety of systems. The key to maintaining high levels of availability in these environments to reduce

the risk of change. Therefore, Change management is not just about applying changes when required, but also about controlling and managing these changes effectively. Changes are usually categorized into 3 different types, based on these controls. A normal change refers to changes that must follow the complete change management process. Normal changes are often categorized according to risk and impact to the organization/business. A normal change will proceed through all steps of the change management process and usually those that are categorised as medium or high risk will be reviewed by the Change Advisory Board (CAB). A standard change is a change to a service or infrastructure for which the approach is pre-authorised by change management, and that has an accepted and established procedure to provide a specific change requirement. Emergency changes are reserved for changes intended to repair an error in an IT service that is impacting the business to a high degree or to protect the organisation from a threat. This process can easily become bureaucratic and administrative. The focus should be on putting the right amount of control on each type of change, and having the tools in place to support it correctly.

Standard changes that have been performed many times before and are fully understood should be automated wherever possible. The SDDC and its software-defined components, such as vSAN and its Storage Policy-Based Management engine, make this a much easier objective to achieve. Tasks and activities that would have previously required manual, tightly controlled changes to the environment can now be automated to avoid the biggest cases of change back-outs, namely manually applied changes. Since vSAN is software-based, it also means fewer change management activities in the physical storage infrastructure, which can mean a smaller change scope and less impact to infrastructure and dependent applications and services.

To embrace the agility and flexibility a software-defined data center-based environment can offer, VMware recommends that a “Pre-approved, standard change” first policy should be created and enforced where possible. A critical success factor for optimizing cloud operations is to utilize pre-approved, standard changes whenever and wherever possible; a simple example being that of applying a pre-defined storage policy to a virtual machine upon creation or modifying the storage availability or performance characteristics on a particular object. Storage policy changes are relatively simple to implement, but you should always consider the number of objects which could be affected by that policy change. A policy change that effects a large number of objects should still be treated as a normal change unless the implications of that change are well known. An example of this would be changing the storage policy of a large number of virtual machines from RAID1 to RAID5/6. In this case a rebuild of the storage components for that object is required, so a large chunk of capacity

may be initially required to perform this change and bring the objects back into a compliant state before the capacity is returned. This could become a significant issue if the migration is under some sort of time constraint, or the buffer capacity is not available to create the required components. Some of these implications are further discussed under the Policy management section. Some other examples of changes and their categories are shown in Table 4.3.

Table 4.3 vSAN change category examples

Action	Type of Change
Create vSAN Cluster	Normal
Enable Deduplication/ Compression	Normal* (VMware recommends this be decided and configured at vSAN cluster creation time)
Enable Encryption	Normal* (VMware recommends this be decided and configured at vSAN cluster creation time)
Hardware extension, hardware upgrade	Normal* (Be sure to always verify firmware and driver versions for compatibility. This may seem obvious but is overlooked more times than we care to think about)
Hardware failure replacement	Normal* (Be sure to verify firmware versions in replacement parts as well as the compatibility of "equivalent" replacement hardware shipped as part of an RMA with the vSAN HCL.)
Software/Firmware Updates	Normal* (Be aware of the implications of putting a host into maintenance mode. Pay close attention to the "what-if" information displayed in the Maintenance Mode UI. It will tell you if objects will become inaccessible as a result of putting the host into maintenance mode.)
Change storage wide configuration setting	Normal
Change availability or performance of an object	Standard* (Consider number of objects and implications)
Apply storage policy to new resource	Standard
Create a new storage policy	Standard
Workload storage migration	Standard

As you can see, the potential impact on change management can have a direct effect on the agility and time to value your business stakeholders experience. Because of the potential impact on business stakeholders, a best practice is to actively monitor the effect of the "normal" changes as they're being made in production during the change window. This also represents another advantage of a blended team model; all team members will have been involved by default and will know to monitor "normal" changes as they're being made. Vigilant monitoring when pre-approved changes take place goes a long way to making for a successful change.

Table 4.4 Sample change management optimization KPIs

Objective	KPI	Description
Positive impact on change management by having a blended team	• Ratio of terminated changes to successful changes	This KPI provides a measure of blended team efficiency in planning and executing changes versus the baseline of the same metric for changes involving siloed teams.
Increase in number of automated Standard Changes	• Total number of automated changes versus manual changes per month	This KPI provides a measure of the number of automated changes being completed which indirectly reflects a reduction in the cost of operations.
Increase number of pre-approved Standard Changes	• Total number of pre-approved Standard Changes compared to Normal and Emergency Changes	This KPI provides a measure of what should be an upward trend in the number of Standard Changes with a downward trend in the number of Emergency and Normal Changes.

Automated provisioning

Many organizations have focused on being able to automatically provision virtual infrastructure and workloads. This may be in the form of offering Infrastructure as a Service to users via an API or self-service consumption model, or simply making it much faster for IT to deploy virtual infrastructure and workloads on the users' behalf.

With vSAN and Storage Policy-Based Management, storage policies can be integrated into the automated provisioning of the virtual infrastructure or workloads being deployed. When these capabilities are used to augment other areas of the SDDC such as NSX, true policy-based automated deployment becomes a reality. This means that workloads can be deployed based on a set of distinct policies that dictate where the workloads are placed, how the workload data should be stored, and with what performance and/or resiliency characteristics, as well as how it may be isolated and protected from a networking and security perspective. This capability speeds up the entire provisioning process, resulting in shorter wait times and removing the probability of human error in both placement and configuration. With the software-defined nature of the SDDC, IT now can go that "last mile" in fully automating the virtual infrastructure or workload provisioning. This concept reinforces the need for blended teams, as described earlier in the book. Success with this level of integration and automation requires virtual compute, network, security, and storage resources to work very closely together.

The vSAN API can be easily leveraged by several Cloud Management Platforms (CMPs) to provide hyper-converged storage services automatically. Some examples of CMPs include vRealize Automation, OpenStack, and VMware® Integrated OpenStack. Given its tight integration of vSAN and SPBM, vRealize Automation is particularly interesting. vRealize Automation is VMware's Cloud Management Platform, and it delivers Day One service provisioning and Day Two operational capabilities across private, public, and hybrid cloud platforms. These capabilities help IT organizations automate core IT processes; speed up infrastructure delivery and get the most out of both equipment and people resources.

vRealize Automation's integration with vSAN enables VM-granular control based on the Storage Policy-Based Management framework. vSphere VM storage policies can be integrated to the vRealize Automation service catalog, allowing you to dynamically assign individual storage policies to virtual machines and virtual disks based on their storage requirements (performance, availability, capacity, etc). This is demonstrated in Figure 4.7.

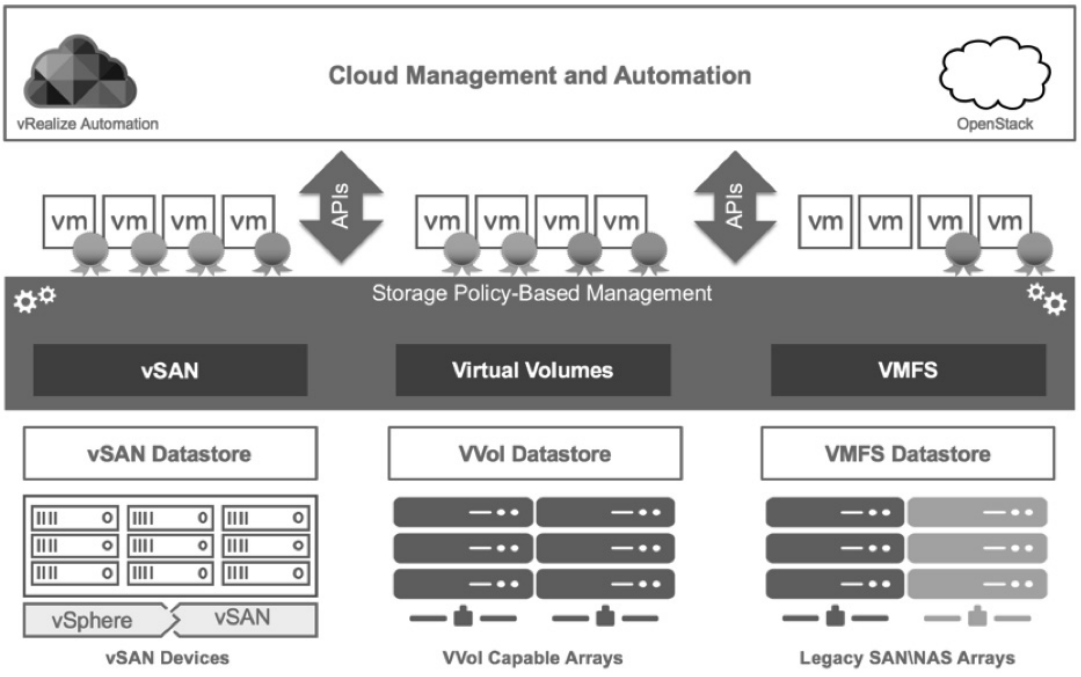

Figure 4.7 vRealize Automation software-defined storage integration

A few examples of how the integration can be used are:

- Applying a SPBM policy when provisioning a VM in vRealize Automation. The policy will be set after VM's have been created (Day 1)
- Applying different policies to VM disks and VM when provisioning a VM in vRealize Automation (Day 1)
- Changing storage policies of existing VM's in vRealize Automation. The policies will be changed. Disks will be migrated by demand (Day 2)
- Migrating a VM to a different vCenter based on a storage policy change in vRealize Automation (Day 2)

This is further demonstrated in Figure 4.8.

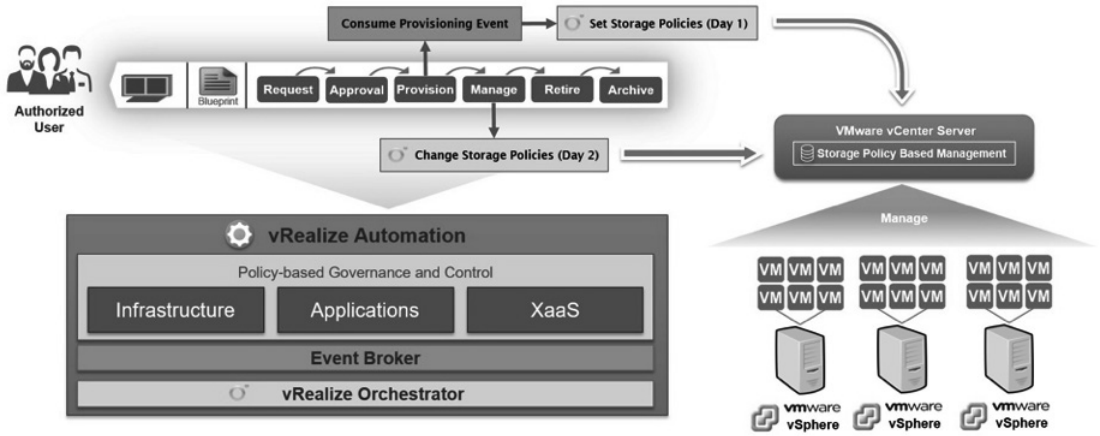

Figure 4.8 vRealize Automation Day 1 and Day 2 software-defined storage operations

Since vRealize Automation makes use of vRealize Orchestrator, you can extend out-of-the-box integrations using the Event Broker Service and XaaS features. The Event Broker Service allows you to subscribe vRealize Orchestrator workflows to specific events, essentially modifying the behavior of the vSAN native integration. XaaS allows you to publish vRealize Orchestrator workflows in vRealize Automation, which can then be directly invoked from the vRealize Automation service catalog.

This tight integration between vRealize Automation and the SDDC components (vSAN, NSX) ultimately allows end users, or IT on behalf of end users, to deploy full application stacks from the vRealize Automation service catalog. This enables end users to quickly deploy a fully configured, performant, secure, redundant and networked application stack, while providing IT the benefit of a standardized and repeatable process. IT has the added benefit of using blueprints which,

since they are text based, can be treated as "infrastructure as code," allowing version control as part of the overall lifecycle management approach for a service such as Infrastructure as a Service.

Table 4.5 Sample integrated vSAN provisioning optimization KPIs

Objective	KPI	Description
Decrease average end-to-end workload provisioning time due to automating and continuously improving storage and continuity related steps.	• End-to-end monthly average workload provisioning time	This KPI should show a decrease in pre-provisioning and post-provisioning due to automating what had been manual storage-related activities.
Increase workload provisioning success rate due to automating storage-related steps.	• Monthly workload provisioning success rate • Monthly cost of provisioning failure	• Monthly workload provisioning success rate tracks total number of workloads requested, total number of workloads successfully provisioned, and total number of workload provisioning failures month over month which should reflect an increasing trend in the success rate. • Monthly cost of provisioning failure tracks (the estimated or actual cost of workload provisioning failures which should decrease over time due to fewer failures)

Continuity management

As businesses become more and more reliant on a complex network of technology to deliver cutting edge consumer focused services and products, it is important to ensure that these services remain highly available if something unforeseen, such as a natural disaster, were to strike. Continuity management aims to manage and plan for these unforeseen risks that could seriously impact IT's ability to deliver these services, and be able to successfully recover these services if something were to happen.

Planning plays a major part in any form of continuity management. A clear and thorough understanding of the services that are provided, and the value of those services to the business, helps establish suitable recovery point objectives (RPO) and recovery time objectives (RTO). Once these are established, an architecture and set of policies can be defined and put in place to ensure we adhere to these set objectives.

vSAN provides a number of deployment options as well as a number of features which can help avoid outages or allow recovery in a relatively short period of time. vSAN is a distributed storage system, but you must take care when defining exactly how and where its data objects are stored. One of the major benefits of vSAN, as it relates to continuity management, is cost. The cost of a disaster recovery site can be prohibitive for many organizations. As a result, a number of businesses have inadequate or no disaster recovery plans at all, which introduces considerable risk. One of the more significant costs of a disaster recovery site is the IT systems infrastructure, including server hardware, storage, and replication software. vSAN allows you to deploy inexpensive industry-standard x86 server components in place of large, upfront investments. Because direct attached drives are used to create the shared storage, there is no dependency on external shared storage hardware. This helps reduce the total cost of the solution, while providing sufficient capacity, reliability, and performance. Software-based replication can provide asynchronous virtual machine replication with low RPO's. Replication can be configured on a per-virtual machine basis, enabling precise control over which workloads are protected.

Carefully consider exactly how vSAN data objects are stored, when looking at redundancy and continuity. vSAN offers several levels of redundancy control, as shown in Figure 4.9. The first level of redundancy control, which can be defined in vSAN is number of failures to tolerate (FTT). This is defined at a policy level and can be applied to virtual machine objects. This policy setting defines the number of host and device failures that a virtual machine object can tolerate. For n failures tolerated, each piece of data written is stored in n+1 places, including parity copies if using RAID 5 or RAID 6. The default value for FTT is 1 and the maximum value is 3 for RAID 1 failure tolerance method, and 2 for RAID5/6. vSAN requires a minimum number of hosts, depending on the failure tolerance method and number of failures to tolerate (FTT) configuration. For example, a minimum of three hosts is needed for FTT=1 with RAID-1 mirroring. A minimum of four hosts is required for FTT=1 with RAID-5 erasure coding.

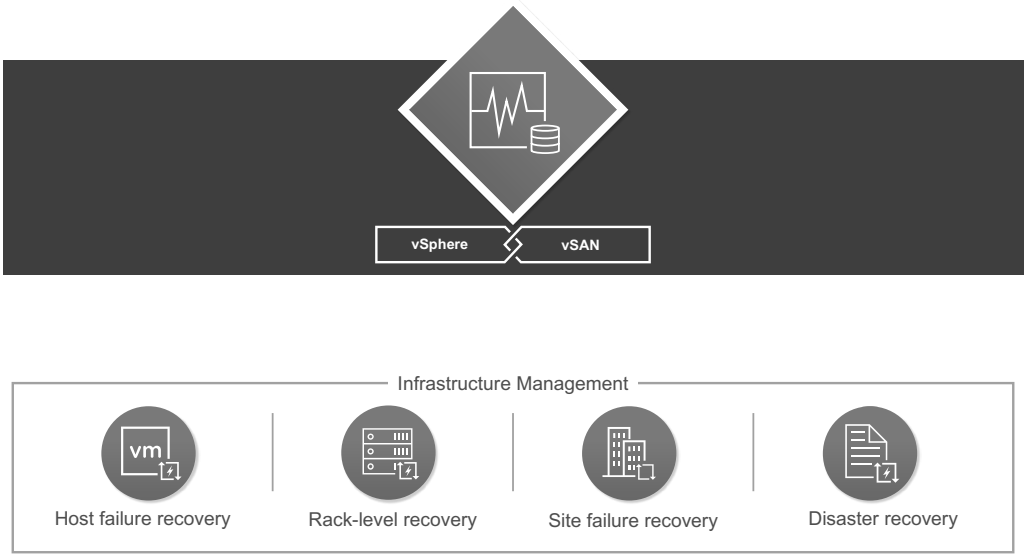

Figure 4.9 vSAN redundancy control options

"Fault domain" is a term that comes up often in availability discussions. In IT, a fault domain usually refers to a group of servers, storage, and/or networking components that would be impacted collectively by an outage. A common example of this is a server rack. If a top-of-rack switch or the power distribution unit for a server rack fails, it would take all the servers in that rack offline, even though the server hardware is functioning properly. That server rack is considered a fault domain. Each host in a vSAN cluster is an implicit fault domain. vSAN automatically distributes components of a vSAN object across fault domains in a cluster, based on the FTT rule in the assigned storage policy. The failure of a disk or entire host can be tolerated in a default deployment. However, this does not protect against the failure of larger fault domains such as an entire server rack. It is possible that multiple components that make up an object could reside in the same server rack. If there is a rack failure, the object would be offline. To mitigate this risk, servers in a vSAN cluster should be placed across server racks and a fault domain should be configured for each rack. This is the second level of redundancy control. Configuring fault domains in this way instructs vSAN to distribute components across server racks to eliminate the risk of a rack failure taking multiple objects offline. This feature is commonly referred to as "Rack Awareness".

The third level of redundancy control is provided by deploying vSAN as a stretched cluster. vSAN stretched clusters are created across two sites and provide resiliency against the loss of an entire site. The two sites could be separate rooms at opposite ends of the same building, two buildings on the same campus, two campuses in separate cities, and so on. There are many possibilities, but always bear in mind the RPO, RTO and recovery strategy objectives. There are also minimum requirements of a 10Gbps network connection with 5ms round trip time latency

between sites. In addition to the hosts at each site, a vSAN witness host is deployed to a third site, as shown in Figure 4.10. Its purpose is to enable the cluster to achieve quorum when one of the two main data sites is offline. The witness host also acts as "tie-breaker" in split-brain situations.

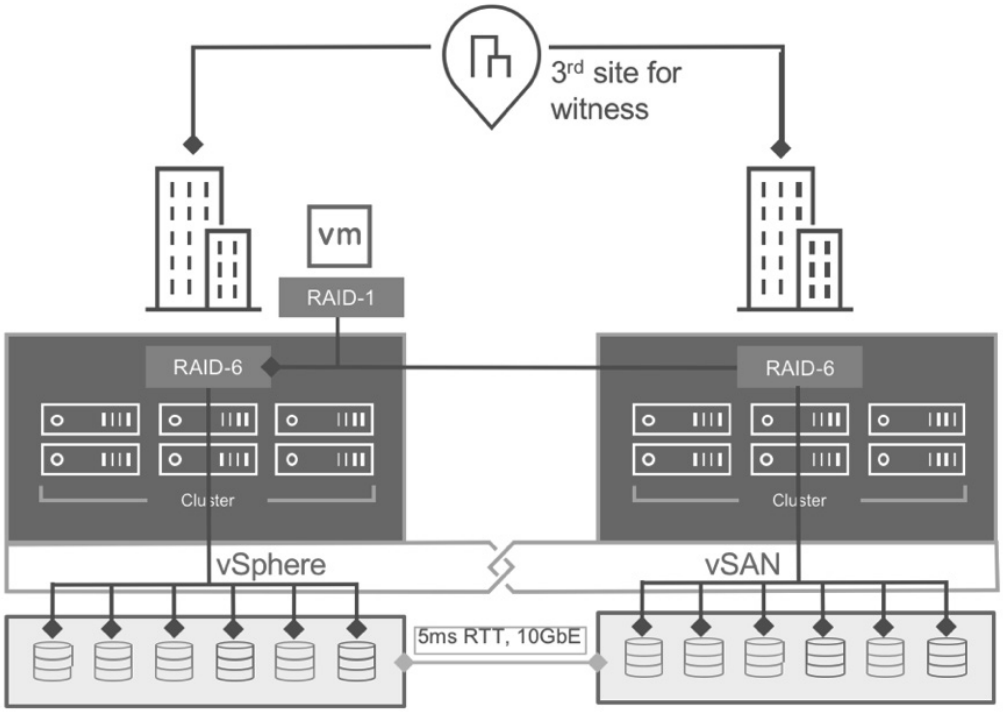

Figure 4.10 Stretched vSAN cluster deployment

vSAN is integrated tightly with vSphere High Availability (HA). If a site goes offline unexpectedly, vSphere HA can automatically restart the virtual machines affected by the outage at the other site. There is no data loss because data is synchronously written to each site. The virtual machines will begin the restart process in a matter of seconds, helping you meet your RTO and minimize downtime.

vSAN stretched clusters are also beneficial in planned downtime and disaster avoidance situations. Virtual machines at one site can be migrated to the other site with VMware vMotion. Issues such as an impending storm or rising flood waters typically provide at least some time to prepare before disaster strikes. Virtual machines can easily be migrated out of harm's way in a vSAN stretched cluster environment.

vSAN has a number of options to help you achieve even the most stringent of RTO and RPO's. However, it doesn't replace the need for a well-defined continuity plan, in case something were to happen. A continuity plan provides a central point of reference of what to do and how to recover business critical services quickly and effectively. It's important to document the plan, because once the recovery steps are understood they can potentially be automated using a site recovery tool such as VMware Site Recovery Manager™ (SRM).

SRM includes the ability to precisely control the startup order of virtual machines, and it automates IP address changes when virtual machines fail over. Testing recovery plans and steps with SRM is non-disruptive, which enables frequent testing. Frequent testing leads to higher levels of confidence that recovery will work as planned. History reports are generated with every test and failover event, providing documentation to satisfy organization and regulatory requirements. The blended team model discussed previously also plays a major role here. IT services are backed by compute, storage and network resources, and the blended team model can help recover or avoid outages by being able to use a central pool of skilled resources working together to achieve the necessary continuity requirements and goals.

Table 4.6 Sample continuity management KPIs

Objective	KPI	Description
Decrease number of outages which are not resolved by the blended team within the recovery objectives	• % of outages resolved within RTO meeting RPO objectives	This KPI provides a measure of how effective the blended team is in resolving outages.
Positive impact on continuity management by having a blended team	• Ratio of recovery times using blended and siloed teams	This KPI provides a measure of blended team efficiency in planning and executing recovery tasks versus the baseline of the same metric for tasks involving siloed teams.

Data security management

Organisations today face an ever-increasing list of statutory and regulatory compliance obligations. The majority of these are positioned to combat issues in areas such as Data Governance, Data Protection, Operational Risk and Information Security. Organisations now need to comply with a wide range of information-related regulations such as HIPAA, PCI DSS and GDPR. They also need to deal with an increasing exposure to rapidly mutating, sophisticated threats to their information and information assets. These threats exploit a diversity of technical vulnerabilities in IT systems, as well as loopholes in procedures and the behavioural characteristics of employees.

How can we address some of these issues and reduce the risk to our organisation? Previously we discussed the need for a blended team model and the consolidation of skill sets in order to deliver and operate IT services effectively. Information security is one of the skillsets that almost definitely needs to be included in this team. Information security and risk management should be addressed as early as possible in the cloud service lifecycle, namely during what we call the Service Definition process.

The Service Definition process activities focus on identifying and defining key service elements required to manage and deliver a true cloud service effectively and efficiently. The Service Definition process applies a 360° view to defining a service and prescribes a collaborative approach by including all associated stakeholders in the process. This 360° view can drive collaboration with business, financial, security, operational, infrastructure, and application developer stakeholders. Following this method allows for a risk analysis and countermeasure identification exercise to be done for the service being defined, as well as any related functional perspectives, for example cloud infrastructure changes or explicit regulatory compliance requirements. These data requirements and countermeasures can then be considered when the service is developed, tested and finally provided to consumers. If a service had distinct data confidentiality or sovereignty requirements, these could then be provided from an infrastructure perspective. This could involve architecting specific clusters to handle these requirements. For example building a vSAN cluster in a specific country that can host sensitive workloads and subsequently enable data encryption either at rest using vSAN encryption or in transit using VM encryption or a combination of both.

Data at Rest Encryption requirements can be satisfied with vSAN Encryption. vSAN datastore encryption uses an AES 256 cipher and eliminates the extra cost, limitations, and complexity associated with purchasing and maintaining self-encrypting drives. vSAN datastore encryption is enabled and configured at the datastore level. When this feature is enabled, every object on the vSAN datastore is encrypted. This is different to standard vSphere VM encryption which can be applied at a VM level. In vSAN, data is encrypted when it is written to persistent media in the cache and capacity tiers of a vSAN datastore. Encryption occurs just above the device driver layer of the vSphere storage stack, which means it is compatible with all vSAN features, such as deduplication and compression, RAID-5/6 erasure coding, and stretched cluster configurations. Because vSAN encrypts data after dedupe and compression, storage efficiencies are preserved, as opposed to standard vSphere VM encryption. In standard vSphere VM encryption, vSAN receives an encrypted data stream and hence offers minimal storage efficiency. All vSphere features including VMware vSphere® vMotion ®, VMware vSphere® Distributed™ Resource Scheduler (DRS), VMware vSphere® High Availability (vSphere HA), and VMware vSphere® Replication® are supported with vSAN encryption. A Key Management Server (KMS) is required to enable and use vSAN encryption. Ensure the KMS infrastructure has high availability and keep it separate from the vSAN infrastructure to avoid circular dependencies. Multiple KMS vendors are compatible, including HyTrust, Gemalto (SafeNet), Thales e-Security (Vormetric), and CloudLink® as shown in Figure 4.11.

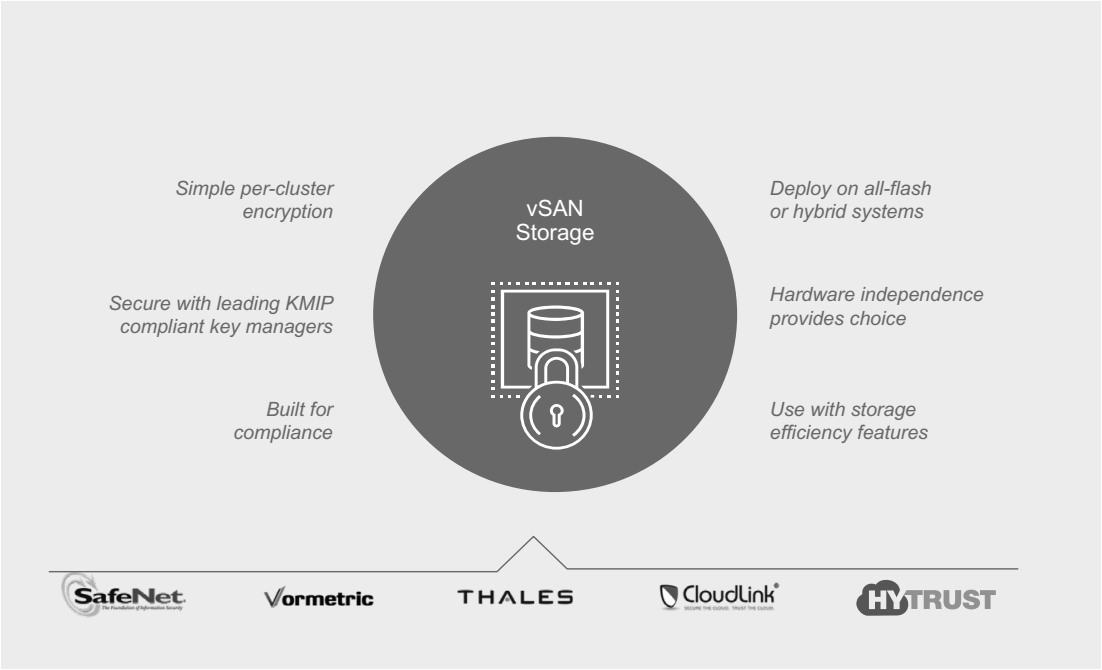

Figure 4.11 Third party key management server providers

Table 4.7 Sample Data Security Management KPIs

Objective	KPI	Description
Increase security posture in newly developed services	• % increase in SLA conformance to security clauses	This provides a measure of how effective the service definition process and the blended team is in involving all required stakeholders during the service lifecycle.
Build security in. Decrease the impact of security incidents against newly developed services	• % decrease in the impact of security incidents	This KPI provides a measure of how effective the blended team is in developing services that mitigate risks and exposure.

Policy management

Traditional storage solutions commonly use LUNs or volumes. A LUN or a volume is configured with a specific disk configuration, such as RAID, to provide a specific level of performance and availability. The challenge with this model is each LUN or volume is confined to providing only one level of service, regardless of the workloads that it contains. This leads to provisioning numerous LUNs or volumes to provide the right levels of storage services for various workload requirements. Maintaining many LUNs or volumes increases complexity. Deployment and management of workloads and storage in traditional storage environments are often a manual process that is time-consuming and error prone.

Storage Policy-Based Management aims to address this complexity and lack of flexibility by enabling precise control of storage services. A SPBM defined policy aims to meet different consumer driven application service level requirements by providing the right storage requirements and characteristics for any particular application or service via their underlying virtual machines and/or virtual disks, an example of which is shown in Figure 4.12. A policy essentially looks to place data in an intelligent fashion to cater for availability, capacity and performance requirements, or required levels of service.

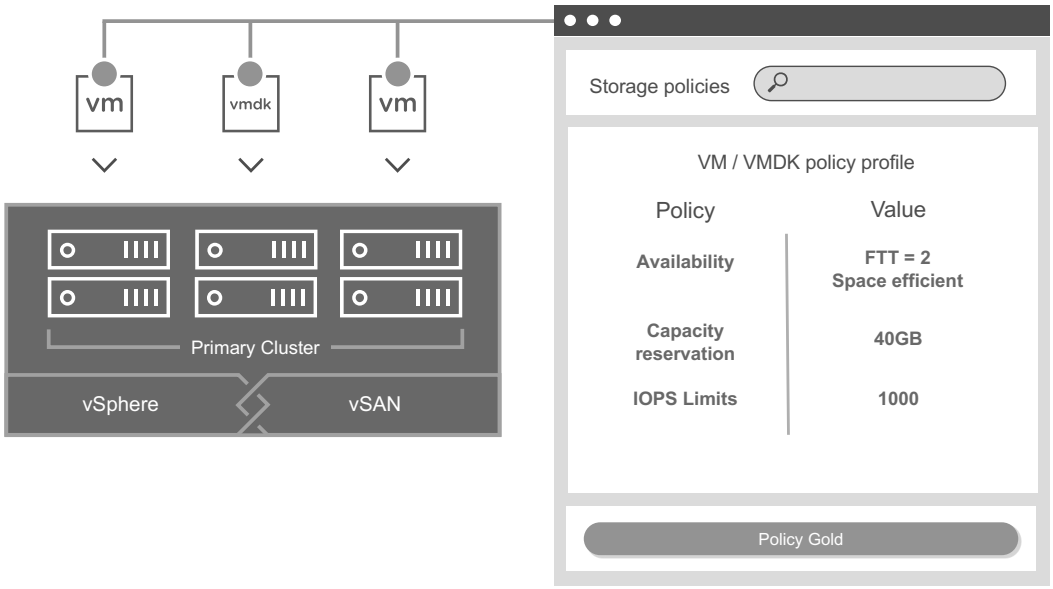

Figure 4.12 vSAN storage policy example

Storage policies are created and managed using the vSphere® Web Client (or increasingly its replacement, the HTML5-based vSphere® Client™), as well as the vSAN APIs and PowerCLI cmdlets. These policies can be assigned to virtual machines and individual objects such as a virtual disk. Storage policies are easily changed or reassigned if application requirements change. These modifications are performed with no downtime and without the need to migrate virtual machines from one datastore to another. SPBM makes it possible to assign and modify service levels with precision on a per-virtual machine or virtual disk basis. vSAN provides a number of storage capabilities which can be configured within a policy as shown in Table 4.8.

Table 4.8 vSAN storage capabilities

Storage Capability	Action
Site Disaster tolerance (PFTT)	Defines the number of host and device failures that a virtual machine object can tolerate. For n failures tolerated, each piece of data written is stored in n+1 places, including parity blocks if using RAID 5 or RAID 6.
Failures to tolerate (SFTT)	In a stretched cluster, this rule defines the number of additional host failures that the object can tolerate after the number of site failures defined by PFTT is reached. For example, if PFTT = 1 and SFTT = 2, and one site is unavailable, then the cluster can tolerate two additional host failures.
Affinity	In a stretched cluster, this rule is available only if the Primary level of failures to tolerate is set to 0. You can set the Affinity rule to None, Preferred, or Secondary. This rule enables you to limit virtual machine objects to a selected site in the stretched cluster.
Number of disk stripes per object	The minimum number of capacity devices across which each replica of a virtual machine object is striped. A value higher than 1 might result in better performance, but also results in higher use of system resources.
Failure tolerance method	Specifies whether the data replication method optimizes for Performance, i.e., RAID-1 (Mirroring), or Capacity, i.e., RAID-5/6 (Erasure Coding).
Object space reservation	Percentage of the logical size of the virtual machine disk (vmdk) object that must be reserved when deploying virtual machines.
Flash read cache reservation	Flash capacity reserved as read cache for the virtual machine object. Specified as a percentage of the logical size of the virtual machine disk (vmdk) object. Reserved flash capacity cannot be used by other objects. Unreserved flash is shared fairly among all objects. This option should be used only to address specific performance issues and is primarily relevant in a hybrid environment.
Force provisioning	If the option is set to Yes, the object is provisioned even if the Number of failures to tolerate, Number of disk stripes per object, and Flash read cache reservation policies specified in the storage policy is not satisfiable by the datastore. Use this parameter in bootstrapping scenarios and during an outage when standard provisioning is no longer possible.
Disable object checksum	If the option is set to No, the object calculates checksum information to ensure the integrity of its data. If this option is set to Yes, the object does not calculate checksum information.
IOPS limit for object	Defines the IOPS limit for an object, such as a VMDK. IOPS is calculated as the number of I/O operations, using a weighted size.

When defining a policy, vSAN has “what if” APIs so it can show what the result would be of having such a policy applied to a VM of a certain size. This is extremely useful as it can give you an idea of what the “cost” is of certain selected capabilities and decisions.

Storage policy changes themselves are relatively simple to implement, but carefully consider when you are looking to change a policy on a set of objects. The decision made around different policy changes could mean that a rebuild of an object or set of objects is required. This could mean an impact on both time and more than likely storage capacity, as an object would need to be created again as part of a rebuild before its older copy is removed and the object is returned to its new compliance state.

Table 4.9 Policy changes requiring a rebuild

Policy Change	Rebuild Required?	Comment
Increasing/Decreasing Number of Failures To Tolerate	No	As long as (a) RAID protection is unchanged and (b) Read Cache Reservation = 0 (hybrid)
Enabling/Disabling checksum	No*	*Rebuild does not take place if checksum is disabled on an object. If checksum is enabled on an object where it is currently disabled then a rebuild does occur
Increasing/Decreasing Stripe Width	Yes	
Changing RAID Protection	Yes	RAID-1 to/from RAID-5/6, and vice-versa. RAID-5 to/from RAID-6, and vice-versa.
Increasing the Object Space Reservation	Yes*	*Rebuild only takes place when the Object Space Reservation of the object is 0, and a new policy is applied that contains an Object Space Reservation value greater than 0. If the Object Space Reservation is already greater than 0 and it is increased no rebuild takes place
Increasing the Read Cache Reservation	Yes	Applies to hybrid only

In summary, policy changes can easily be made on-the-fly to change your VM's storage requirements, should your application requirements need it. The flexibility provided by changing storage policies on the fly overcomes some of the traditional operational issues, such as moving a VM to a new datastore to attempt to gain more performance/ resources, and then moving the VM back because the LUN/NFS datastore did not meet the expected results. Changing policies on-the-fly also alleviates the significant number of man hours and the specialized knowledge required for the planning, change control, implementation, and validation associated with making traditional storage changes. At the same time, understand that changing policies on-the-fly in this way may require additional capacity on the vSAN datastore and can also impact vSAN performance, since changes like this can instantiate a significant amount of rebuild/resync traffic on the vSAN network. This is especially true if you change a policy that impacts many VMs at the same time.

Table 4.10 Sample Storage Policy-Based Management KPIs

Objective	KPI	Description
Increase customer satisfaction by decreasing downtime	• Difference in time spent moving a customer's workload between service tiers in a physical storage-based environment versus vSAN	This KPI provides a measure of efficiency in day 2 operations which can be translated to OPEX savings.
Increase customer satisfaction by rapidly offering new service tiers to meet changing business needs	• Difference in time spent providing a new service tier to customers in a physical storage-based environment versus vSAN	This KPI provides a measure of efficiency in meeting customer requests and can also be translated into OPEX savings.

Service level management

With IT organizations increasingly moving to the model of becoming internal service providers that supply business-enabling solutions, it is becoming increasingly important for them to be able to define, agree, and adhere to Service Level Agreements (SLAs) with their end consumers. Defining service levels around key criteria such as the Availability and Performance of delivered services, as well as being able to effectively deliver and report on them, is a key premise to maintaining a healthy relationship between IT and the lines of business.

A broader range of IT services is now also expected to ensure the best fit and the most cost-effective use of IT infrastructure. Often IT provides a tiered set of services to provide a level of choice for its consumers. Performance and Availability are two of the key areas which are often used to distinguish between different tier offerings. From an architecture perspective, this is often accomplished by utilizing principles such as availability zones, cluster configuration options such as vSphere HA or DRS, or perhaps deploying workloads on dedicated hardware with distinct characteristics. Traditionally as related to storage, availability tiers would be linked to location or a RAID configuration of a SAN or NAS, for example, which would most likely have to be pre-defined at LUN creation or array deployment time. An example of a traditional tiered storage policy taking into account availability is shown in Table 4.11.

Table 4.11 Traditional availability tiers

Tier	Implementation Characteristics	Relative Cost
Gold	RAID 5/6 (Distributed Parity)	**
Silver	RAID 1 (Mirror)	*
Bronze	Can possibly be made available to customers	***

Also, performance would be most likely be linked to a particular physical disk type such as Magnetic or Flash or entire array type. An example of a traditional tiered storage policy taking into account performance is shown in Table 4.12.

Table 4.12 Traditional performance tiers

Tier	Performance	Relative Cost
Gold	All Flash	***
Silver	Magnetic 10,000 RPM	**
Bronze	Magnetic 7200 RPM	*

A combination of the above-mentioned factors can lead to complex infrastructure architecture and deployment, leading to a high level of operational effort and little flexibility in catering for different workload types.

vSAN's Storage Policy-Based Management makes it far easier to align storage services and service levels with consumers application and workload requirements. The policy-based approach succeeds in aligning storage services with application requirements by shifting the traditional storage operational model from a bottoms-up, array-centric approach to a top-down, VM centric model. After storage policies are configured, the storage consumer can choose their desired application or virtual machine, and the policy engine will read the associated storage resource policy and orchestrate the precise provisioning of storage resources that match the application's storage resource requirements. In this way, performance and availability policies can be applied at a more granular level per disk, as opposed to per LUN or specific hardware array. This means that with vSAN, service level changes for particular workload requirements do not require migrations or physical storage changes between arrays or LUNS, and can be changed on the fly with no down time. An example of an application centric tiered storage policy taking into account performance and availability is shown in Table 4.13.

Table 4.13 Application centric service tiers

Tier	Performance	Availability	Relative Cost
Business Critical	3 stripe width. Object space reservation 100%	FTT = 2	**
Standard	1 stripe width. Object space reservation 0%	FTT = 1	*
Custom	Based on customer requirements	Based on customer requirements	***

From a performance SLA perspective, defining that data objects stored with particular stripe width will provide better performance is probably not detailed enough to underpin rigorous SLA conversations. We probably want to provide some real metric performance data to back this statement up when having SLA discussions with the consumers of the platform, or when making decisions on tangible performance SLA's to deliver on. Performance could be affected by a number of different things, such as disk and host vendor selection, vSAN configuration, policy definition, etc. VMware provides a storage performance testing automation tool, HCIBench, which is released as a Fling and integrates into the vSAN performance diagnostics service. This tool is deployed as a virtual appliance and can fully automate the end-to-end process of deploying test VMs, coordinating workload runs, aggregating test results, and collecting necessary data for troubleshooting purposes. The results can be used to identify configuration options that can be changed to increase performance, and also in conjunction with vSAN performance diagnostics provide detailed performance data to show metrics such as IOPS, throughput, and read write latency of your test workloads and clusters as shown in Figure 4.13. This data will better arm you to have these performance SLA discussions with your consumers.

Figure 4.13 vSAN Performance Service – Cluster Overview

Table 4.14 Sample service level management KPIs

Objective	KPI	Description
Increase customer satisfaction in relation to provided IT services	• % reduction in the number and severity of storage related SLA breaches • % increase in customer satisfaction of SLA achievements via service reviews and satisfaction surveys	This KPI provides a measure of the quality and quantity of delivered services utilizing effective storage policies.
Increase efficiencies of blended team in defining SLA's for services	• Decrease in time spent on developing and agreeing appropriate SLA's	This KPI provides a measure of how efficient the cloud blended team is in terms of providing suitable consumer service tiers.

Incident management

Incident management is another process that is impacted in the context of operationalizing vSAN. As you'll see in the Chapter 6 - Tools for Monitoring & Troubleshooting section later in the book, a number of tools and their advanced capabilities can be applied to vSAN. Some, like vRealize Log Insight, have content packs available to explicitly include vSAN log analysis and events while others, like vRealize Operations Manager, can apply their advanced capabilities like dynamic thresholds and predictive analytics to vSAN configured clusters. These tools, when used separately or together, can significantly streamline the troubleshooting and resolution, including automated remediation, of issues from a full-stack perspective including vSAN.

In addition, because vSAN is software-defined with an available REST API, it can be easily integrated with orchestration engines like vRealize Orchestrator, or intelligent operations tools like vRealize Operations Manager. This provides the opportunity to integrate vSAN directly, as well as via monitoring and troubleshooting tools with third-party products such as ITSM tools from ServiceNow or Cherwell. This can be particularly useful in the context of intelligent operations.

As previously described, a goal of intelligent operations is to identify and resolve issues before they become service or application impacting. In order to track the effectiveness of taking an intelligent operations mindset, it's good to track the ratio of issues resolved reactively, via traditional incident management, to those proactively resolved before they became incidents. One way we like to do this for organizations is to add the capability to mark entries in their ITSM tool's incident management module as "Resolved Proactively," and provide the capability in vRealize Operations Manager to invoke a

vRealize Orchestrator workflow that creates an incident ticket with pertinent information, but mark it as "Resolved Proactively." In this way, they can easily track the ratio of incidents resolved proactively to those resolved reactively for their vSAN enabled infrastructure.

Finally, implementing the blended team described earlier also streamlines incident management. A blended team acting collectively with shared knowledge of the vSAN-backed infrastructure in the context of a software-defined data center is able to resolve issues much faster than the traditional way of having to continually hand-off and communicate between teams in a siloed organization. This is especially pertinent and valuable when under the pressure of having to resolve a critical severity incident.

Table 4.15 Sample incident management KPIs

Objective	KPI	Description
Decreased time to troubleshoot and remediate an incident due to blended team	• % of escalations resolved within the blended team • Average time to resolve escalations with the blended team	This KPI provides a measure of how effective the blended team is in resolving incidents.
Increase in the number of issues proactively resolved	• Number of vSAN-related issues detected and resolved before they become service or application impacting incidents • % reduction in vSAN-related incidents month over month	This KPI provides a measure of how effective proactive monitoring is in identifying and resolving issues before they become service or application impacting.

Consuming vSAN

The end goal for operationalizing vSAN is to support its use. Unlike the constraints placed on IT by using monolithic storage arrays and supporting storage infrastructure, vSAN provides a whole new level of flexibility. Traditional storage forced workloads to contend for the same shared storage resources, which was often the root-cause of compromises in performance, quality of service, and capacity. vSAN, on the other hand, provides the flexibility to easily create modular, cluster-aware storage environments, enabling complete storage traffic separation. This allows cost-effective placement and storage characteristic decisions, such as performance and protection, to be based on the needs of the applications or the business instead of what's available in shared storage.

To be effective, an analysis of workload patterns should be undertaken to drive cluster design in the software-defined environment. While this is true regardless of the underlying storage, having the flexibility provided by vSAN to now include cluster-aware storage optimized for the workload patterns supported in that cluster makes the effort that much more worthwhile. Doing so then makes consuming vSAN and taking advantage of its flexibility incredibly simple yet powerful for IT's stakeholders and users, because of storage policies and the ability to associate them with virtual machines (VM) and virtual machine disks (VMDK).

At the simplest level, a default storage policy can be assigned to a cluster designed to host workload patterns with similar performance and resiliency requirements. In this case, any and all workloads placed in that cluster will automatically have the default storage policy applied. This is not much different than what would occur in a traditional shared storage environment. However, what if a cluster is designed to accommodate a workload pattern but groups of workloads in that pattern can benefit from different storage characteristics, like Dev/Test for example? In this case you might have a storage policy that optimizes for performance to accommodate compiling and unit testing, while a second storage policy can be optimized for capacity to accommodate integration testing and staging.

Coupling storage policies with a storage policy-aware automated provisioning tool like VMware vRealize Automation provides additional levels of flexibility and simplicity while optionally providing IT with control. For example, the end-user can be presented with a list of pre-built storage policies appropriate to their environment based on the end-users answer to a question about workload type. For more control, or to accommodate less mature end-users whichever the case may be, a series of questions can be presented that then programmatically selects a pre-defined storage policy to be applied to the workload during provisioning requiring no decision on the part of the end-user. Accommodating a new application type being introduced can now be as simple as defining a new storage policy, if even necessary.

> **NOTE:** For more information on using vSAN to create cluster-aware storage environments, please refer to "Cost Effective Independent Environments using vSAN.pdf" link provided in Table 8.1 in the "Where you can go for more information" section.

Table 5.1 Sample consuming vSAN KPIs

Objective	KPI	Description
Decrease the amount of time it takes to provide business-specific, segregated storage environments	• Average time to stand-up a new segregated vSAN cluster	This KPI provides a measure of agility in meeting business-specific needs that can best be met by providing a segregated storage infrastructure.

Tools for Monitoring and Troubleshooting

Intelligent operations

The most visible and frequently asked about aspects of operationalizing vSAN are monitoring and troubleshooting. Effective monitoring and troubleshooting begins with the tools. Fortunately, there are a solid set of tools surrounding the vSAN technology available to both enable a key aspect of Intelligent Operations, namely proactive monitoring, as well as provide very effective troubleshooting capabilities. We'll categorize the tools as shown in Figure 6.1.

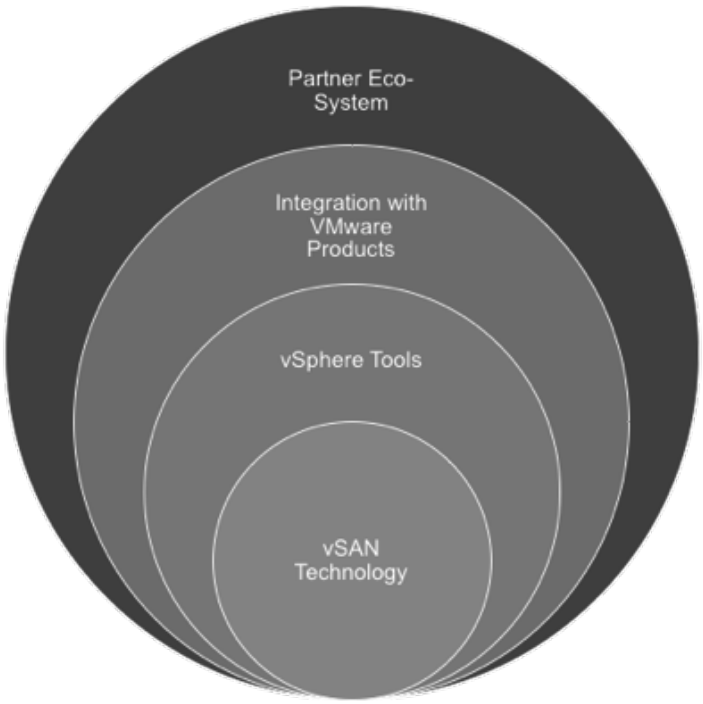

Figure 6.1 vSAN tool ecosystem

These categories include:

- vSphere Tools refers to core vSphere-based tools for vSAN troubleshooting
- Tools provided by VMware Products refers to vSAN monitoring and troubleshooting capabilities provided by the vRealize Operations and vRealize Log Insight
- Partner Eco-System tools refers to additional vSAN troubleshooting tools provided by 3rd party VMware partners

vSphere tools

We'll start with vSphere Tools. vSAN is seamlessly integrated into vSphere and vCenter Server™. As such, a core set of vSAN monitoring and troubleshooting tools are included directly. The primary way of accessing these capabilities are with the vSphere Web Client, vSphere CLI, and vSphere PowerCLI.

vSAN monitoring and troubleshooting tools are accessed via the vSphere Web Client (or the HTML5-based vSphere Client going forward). Using the vSphere Web Client, you can monitor and troubleshoot at the vSAN cluster level as shown in Figure 6.2 and Figure 6.3.

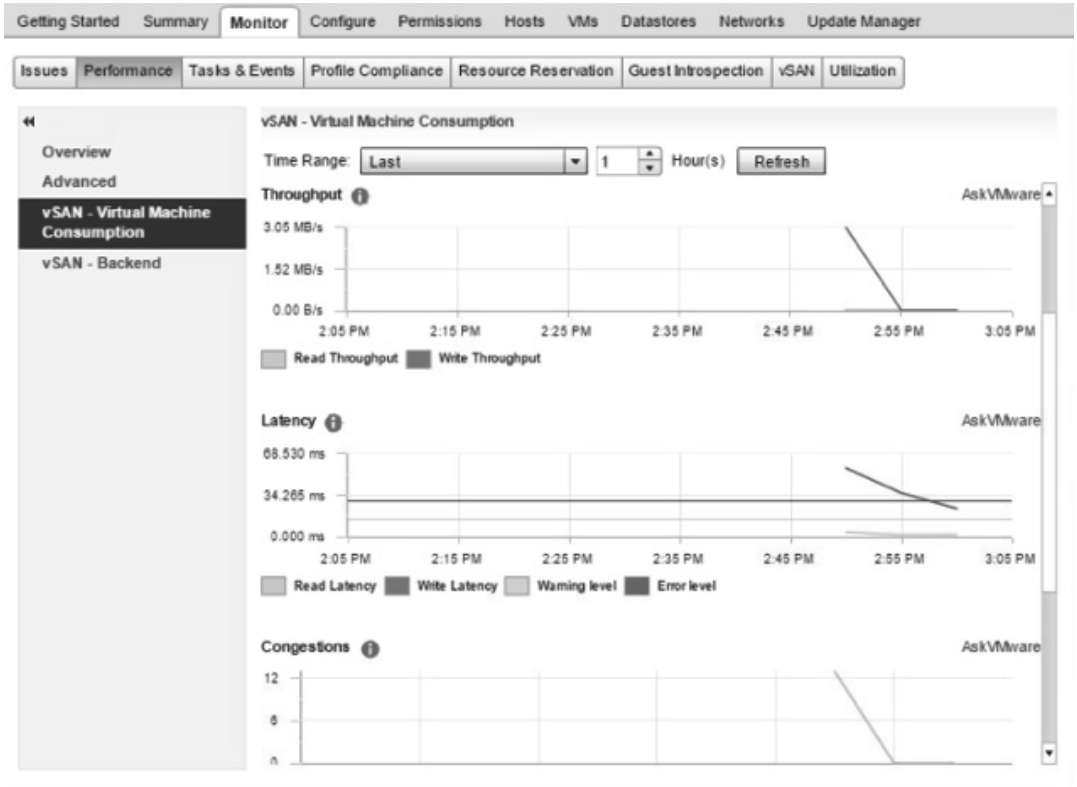

Figure 6.2 vSAN Cluster performance monitoring example

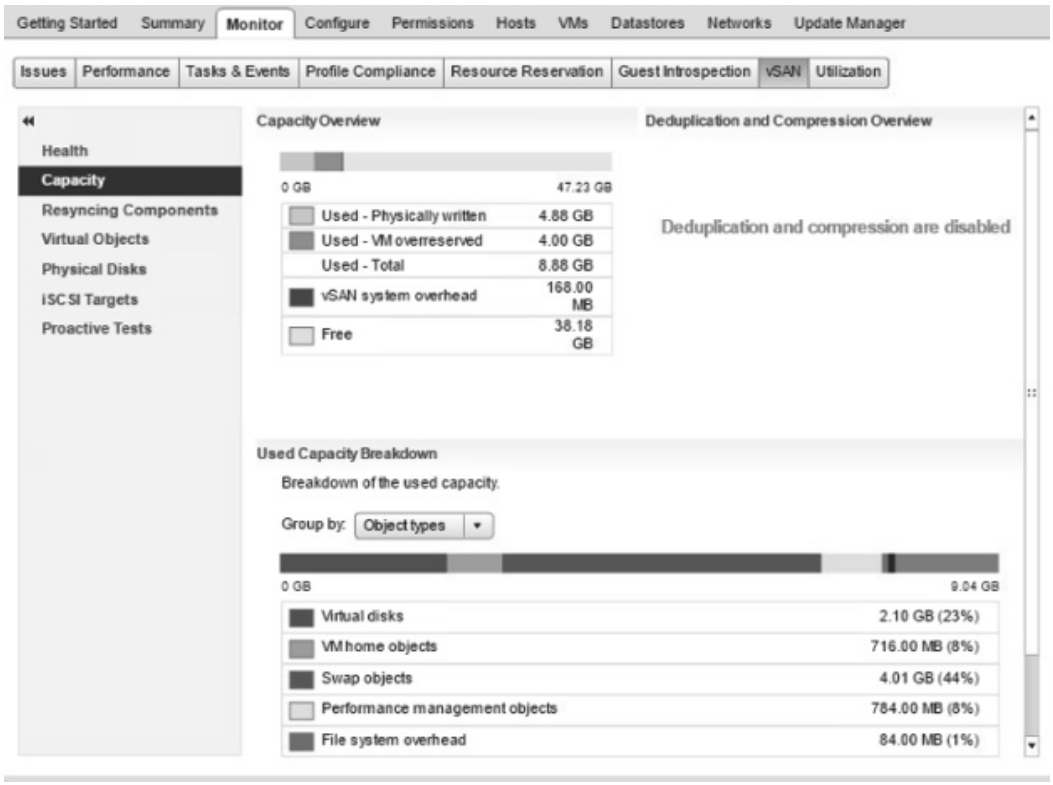

Figure 6.3 vSAN Cluster capacity monitoring example

As well as view overall vSAN health as shown in Figure 6.4.

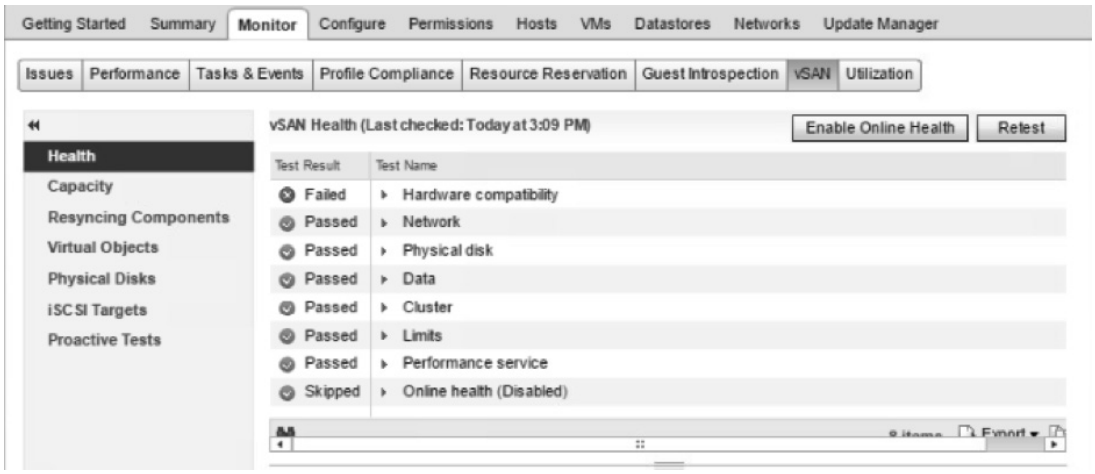

Figure 6.4 vSAN Health Check categories

Table 6.1 contains a brief description of each vSAN Health Check Category.

Table 6.1 vSAN Health Check category descriptions

Health Check Category	Description
Hardware Compatibility	Monitor the cluster components to ensure that they are using supported hardware, software, and drivers.
Network	Monitor vSAN network health
Physical disk	Monitor the health of physical devices in the vSAN cluster.
Data	Monitor vSAN data health.
Cluster	Monitor vSAN cluster health.
Limits	Monitor vSAN cluster limits.
Performance Service	Monitor the health of a vSAN performance service.
Online health	Monitor vSAN cluster health and send to VMware's analytics backend system for advanced analysis. You must participate in the Customer Experience Improvement Program to use online health checks.
vSAN iSCSI target service	Monitor the iSCSI target service, including the network configuration and runtime status.
Encryption	Monitor vSAN encryption health.
Stretched cluster	Monitor the health of a stretched cluster, if applicable.

vSAN periodically retests each health check and updates the results. Health checks can also be run and updated manually. Furthermore, Each Health Check result also includes an 'Ask VMware' button that users can select to open a knowledge base article that describes the Health Check and provides information about how to resolve the issue. Please note that the "Ask VMware" feature requires internet connectivity and that VMware's Customer Experience Improvement Program (CEIP) is turned on.

Each category, in turn, contains more detailed health checks pertinent to that category. An example for vSAN cluster health is show in Figure 6.5.

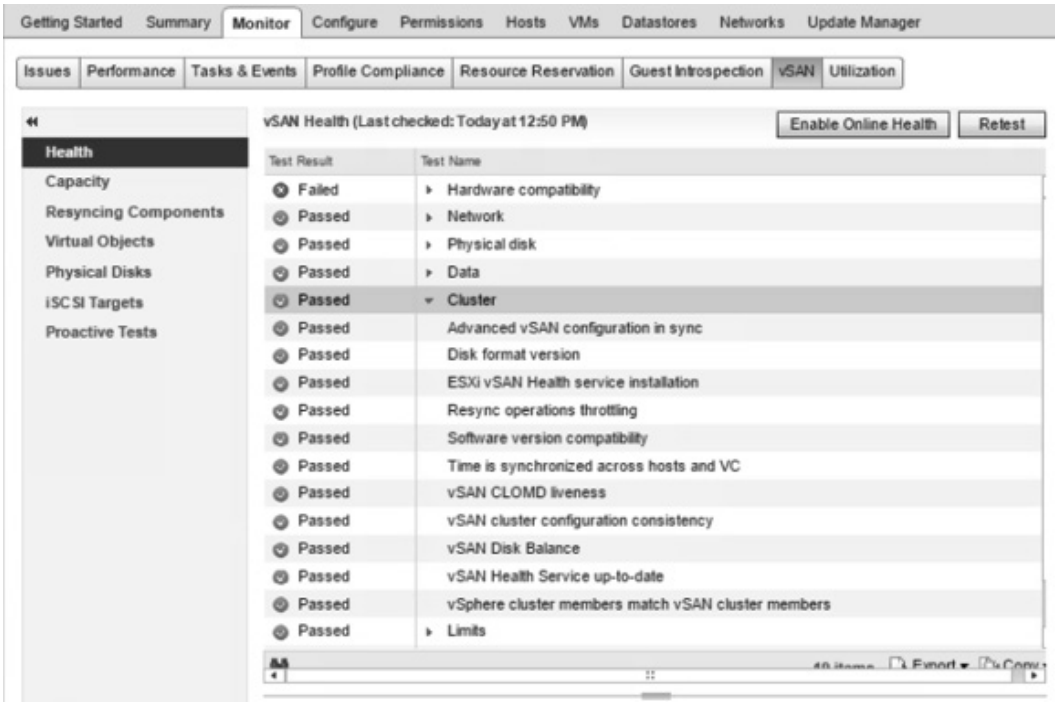

Figure 6.5 vSAN Cluster Health example

You can also monitor and troubleshoot vSAN at the host level using the vSphere Web Client as shown in Figure 6.6.

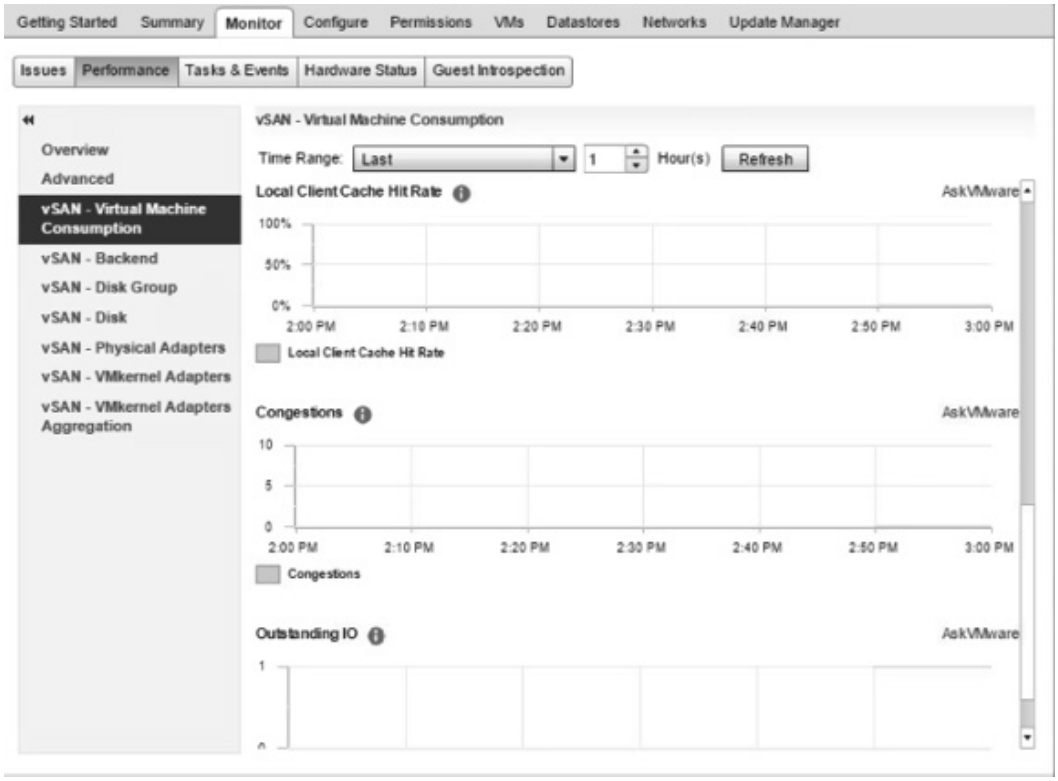

Figure 6.6 Host-based vSAN monitoring example

The command-line interface tools, vSphere, CLI, and VMware PowerCLI are useful for interacting with your vSAN environment especially when vCenter Service is not available.

- vSphere esxcli is a command-line interface that includes vSAN commands to obtain information about vSAN and to troubleshoot a vSAN environment from a host-centric perspective.

- VMware PowerCLI provides a Windows PowerShell interface to the vSphere API. It consists of cmdlets for managing, monitoring, automating, and operating a number of VMware products. In addition to vSAN disk group and stretched cluster-related commands, VMware PowerCLI also provides a number of storage policy related cmdlets.

VMware vRealize intelligent operations tools

Moving up the vSAN monitoring and troubleshooting stack we arrive at Integration with VMware products. There are two VMware products which provide expanded vSAN monitoring and troubleshooting capabilities:

- vRealize Operations

- vRealize Log Insight

vRealize Operations

vRealize Operations provides intelligent operations from application to infrastructure across physical, virtual, and multi-cloud environments. It utilizes predictive analytics and machine learning to effect proactive performance and capacity management which is core to achieving intelligent operations. In the vSAN context, it uses the vRealize Operations Management Pack for vSAN which employs the vCenter and vSAN APIs to collect data and the sophisticated analytics engine to manipulate and analyze the data in ways not possible using the vSphere tools alone. For example, Figure 6.7 shows the vSAN Operations Overview dashboard which in addition to key vSAN storage metrics such as IOPS, throughput, and latency, also provides other metrics that contribute to the health and well-being of the vSAN cluster itself, such as the host count, CPU and Memory utilization, and alert volume.

NOTE: As of vRealize Operations 6.6, the separate management pack is no longer needed as the content it provided has been directly integrated.

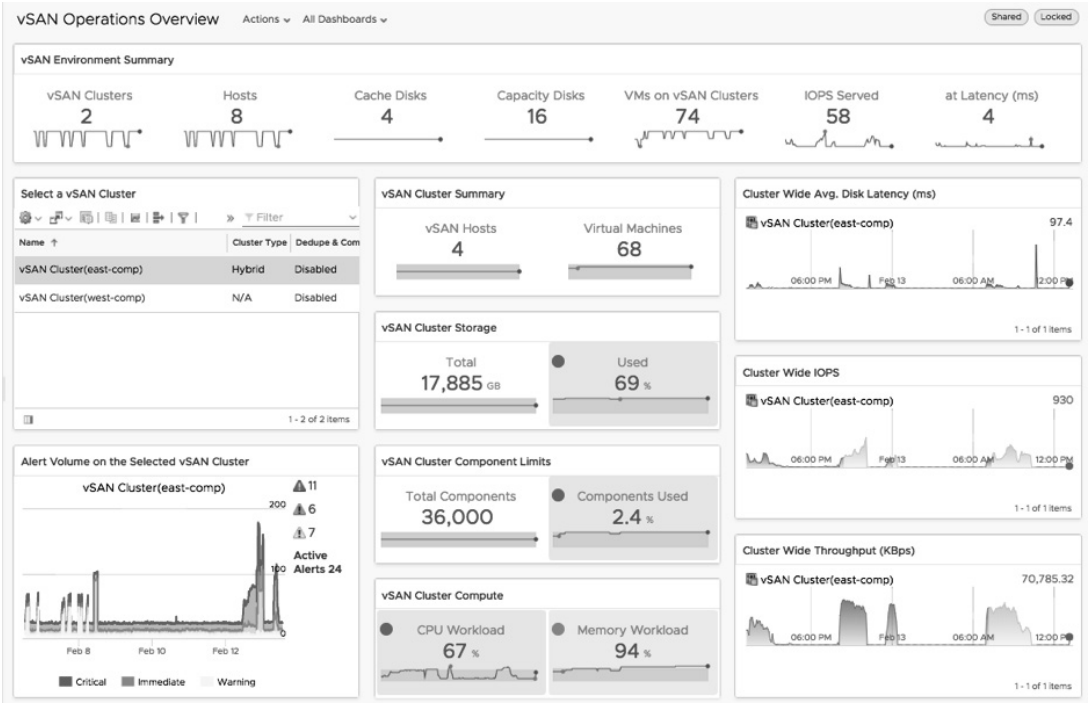

Figure 6.7 vSAN Operations Overview dashboard

In addition to a vSAN operational overview, vRealize Operations provides a vSAN Capacity Overview out-of-the-box as shown in Figure 6.8. This includes a good proactive operations example where, in addition to basic vSAN storage capacity information such as total and remaining, it provides a view into disk usage balance, time remaining before resources are exhausted, and storage reclamation opportunities.

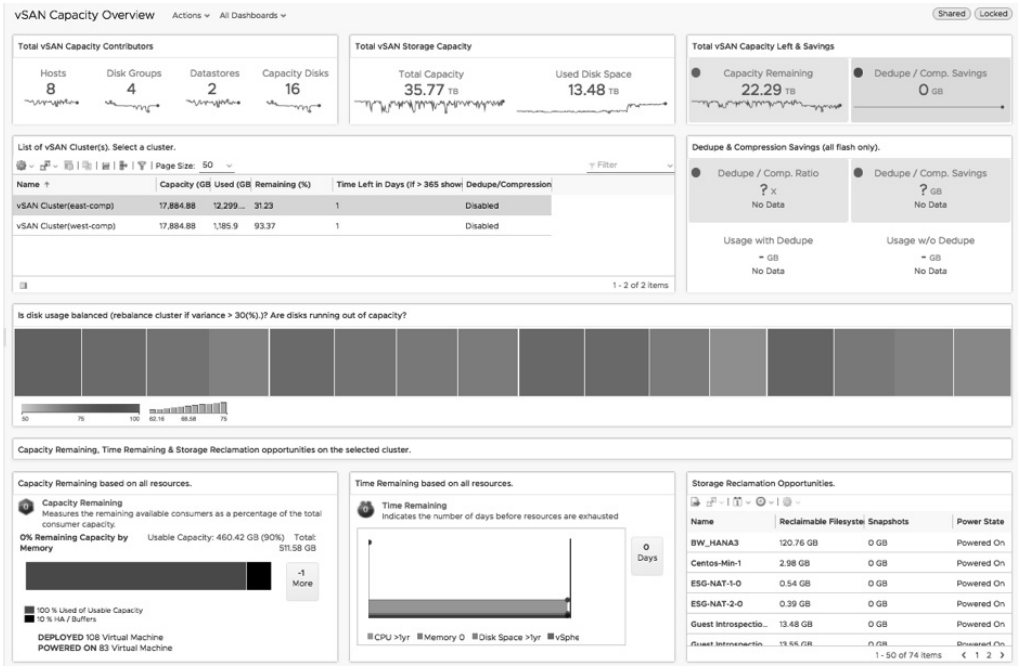

Figure 6.8 vSAN Capacity Overview

One of the critical considerations of operationalizing vSAN is how to troubleshoot issues and, including Intelligent Operations considerations, how to proactively identify, troubleshoot and resolve issues before they become service, application, or end-user impacting. As mentioned previously, Intelligent Operations is very much about culture and mindset change but it also helps to have the right supporting tools. Ideally, such tools can alert you to potential issues, as opposed to ones that are already causing problems, as well as provide context-sensitive troubleshooting capabilities and guided root cause analysis to accelerate resolution. vRealize Operations provides a good example of these troubleshooting capabilities.

The vRealize Operations Management Pack for vSAN provides the data and dashboards needed to support these inherent vRealize Operations capabilities for vSAN specifically. Out-of-the-box it provides a Troubleshoot vSAN dashboard as a great starting point by integrating and displaying data from all of the key vSAN components as shown in Figure 6.9 and Figure 6.10.

Figure 6.9 vRealize Operations' Troubleshoot vSAN dashboard – top

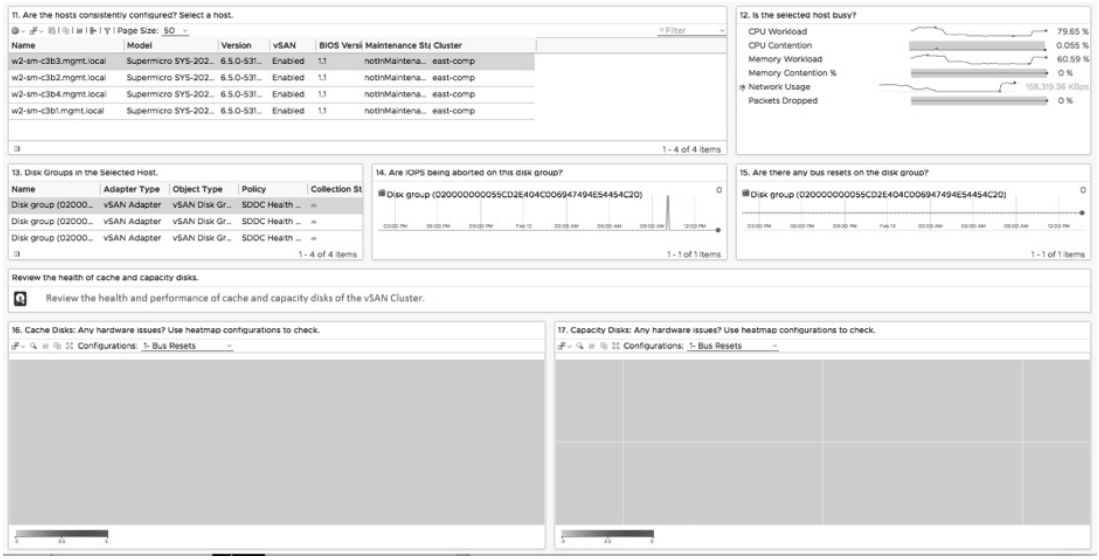

Figure 6.10 vRealize Operations' Troubleshoot vSAN dashboard - bottom

A great example of the context-sensitive troubleshooting and guided root cause analysis capabilities can be shown by following the "More than half of the virtual machines served by this datastore are facing disk write latency" alert shown in Figure 6.11.

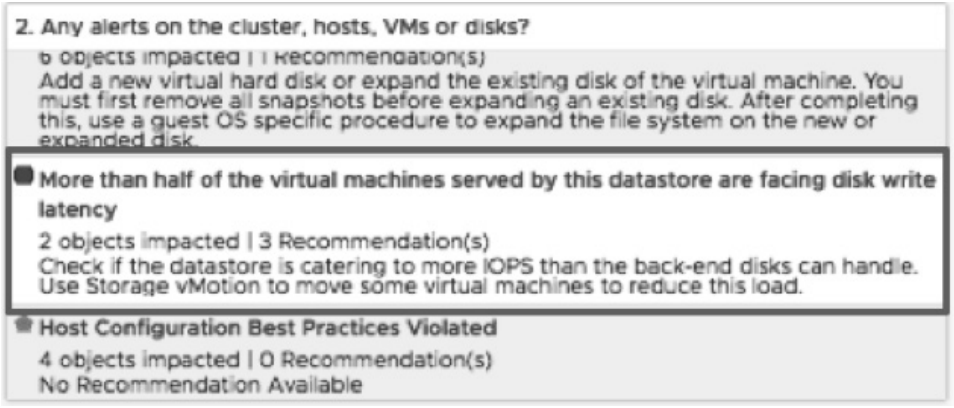

Figure 6.11 Datastore disk write latency health issue

Selecting the "More than half of the virtual machines served by this datastore are facing disk write latency" text presents information about the specific datastores and associated clusters as shown in Figure 6.12.

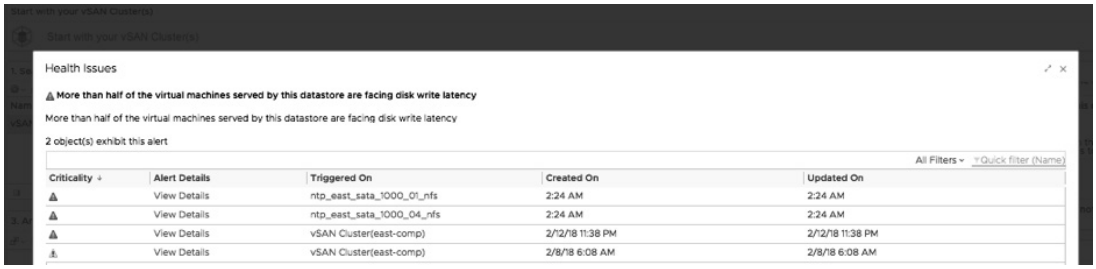

Figure 6.12 Affected datastores and clusters

By selecting a specific component such as the datastore “ntp_east_sata_1000_01_nfs” results in a dashboard displaying detailed information about that datastore as shown in Figure 6.13 as well as a more detailed set of guided root cause analysis recommendations for remediation as shown in Figure 6.14. Note that the datastore complaining about write latency is an NFS target (vs. vSAN) – what’s great about this alert mechanism is that users could choose to Storage vMotion to a better performing vSAN Datastore as a remediation option.

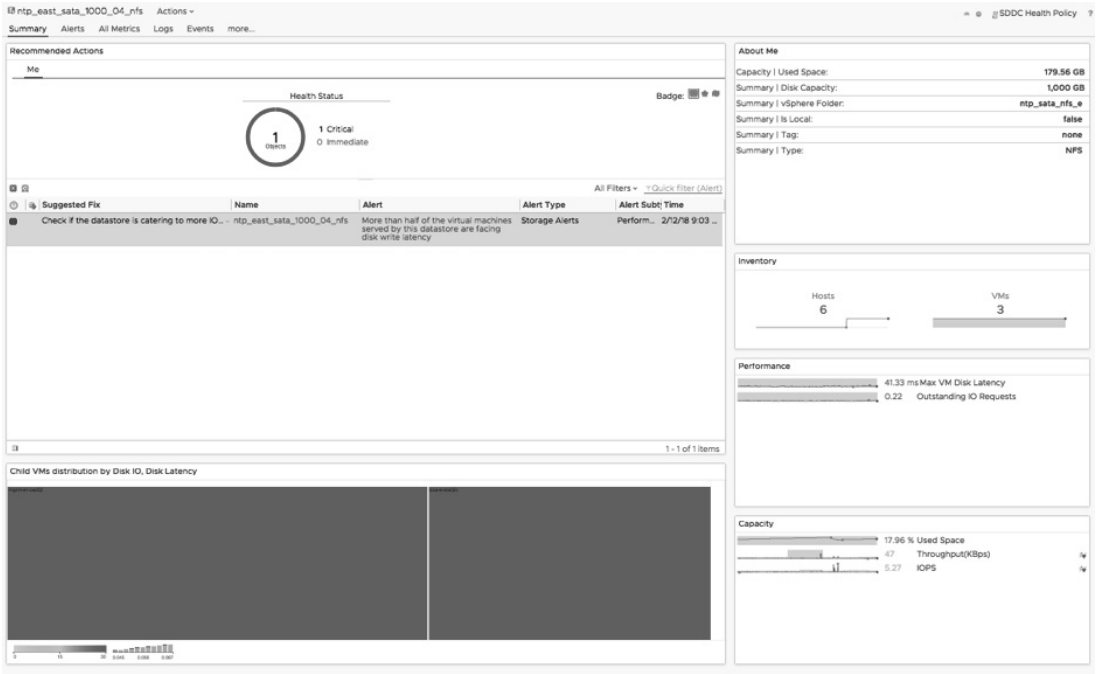

Figure 6.13 Affected vSAN datastore details

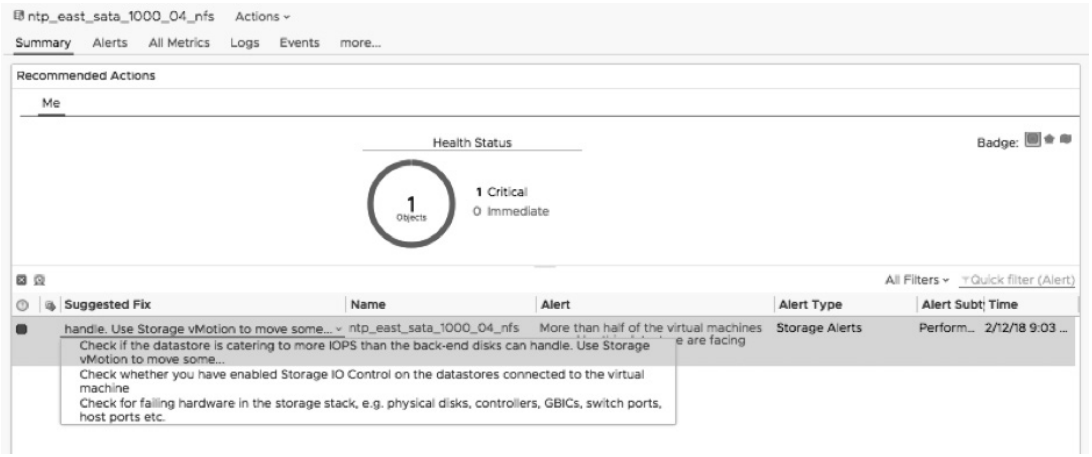

Figure 6.14 Guided root cause analysis recommendations

For last mile analysis, you can also directly access context-sensitive log information in vRealize Log Insight directly while troubleshooting in vRealize Operations as shown in Figure 6.15.

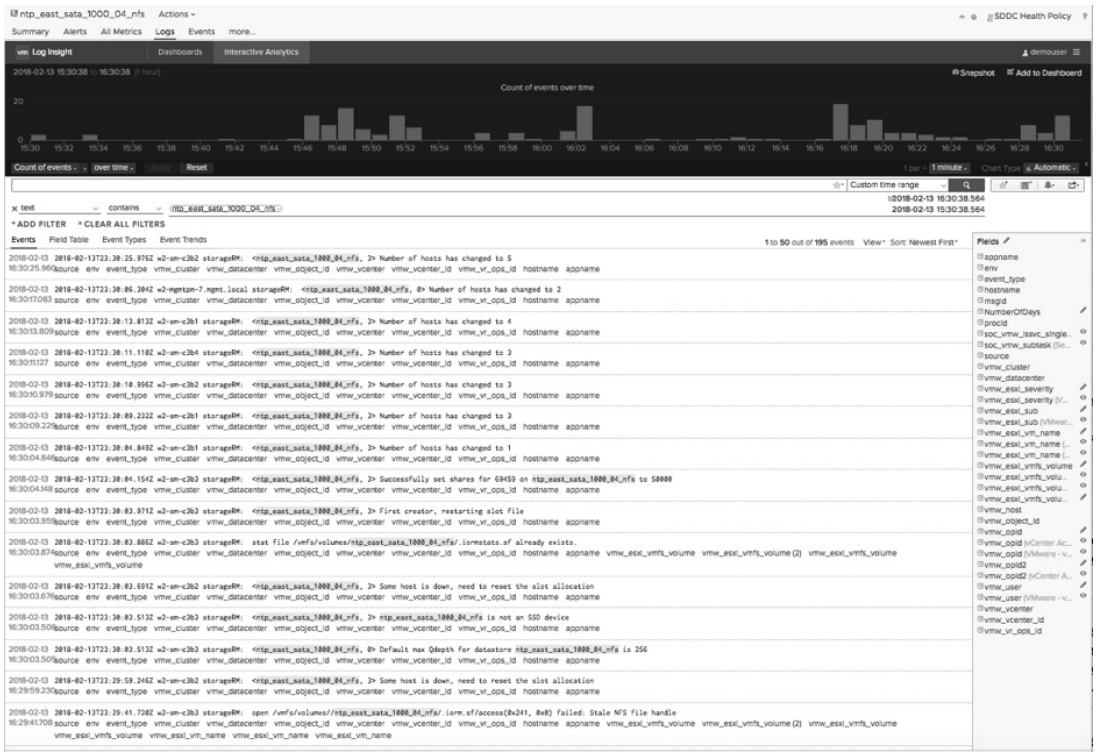

Figure 6.15 Context-sensitive log information

vRealize Log Insight

This brings us to our second VMware Product, vRealize Log Insight. vRealize Log Insight provides highly scalable log management with intuitive, actionable dashboards, sophisticated analytics and broad third-party extensibility. The vRealize Log Insight Content Pack for vSAN provides operational reporting, trending, and alerting visibility for all log data within vSAN. This content pack provides “last mile” troubleshooting to compliment the monitoring and troubleshooting capabilities of vRealize Operations. Out-of-the-box it includes dashboards for all of the most critical vSAN considerations as shown in Table 6.2.

Table 6.2 vRealize Log Insight vSAN dashboards

Dashboard	Description
Host State Information	Overview of how vSAN treats host membership and roles with a primary focus around activities of the host itself, such as additions or changes in host membership to a vSAN cluster
Diskgroup Failures	Collection of diskgroup activity widgets providing visibility into diskgroup failure events
Networking	Displays vSAN network creation and connectivity events
Congestion	Displays events related to how vSAN provides storage traffic flow control
Object Configurations	High-level view of activities related to vSAN objects such as object configuration creation, change, and repair as well as object configuration rebalance, decommissioning, cleanup and vote rebalance
Decommissioning	Displays events related to host maintenance mode as well a disk decommissioning
Configuration Failures	Displays events primarily focused on storage policy-related object configuration activities
Operation Failures	A "failures-based" dashboard focused on vSAN related operations failures such as operations taking too long, component creation failures, and resync operations
Health	Provides an overview of health status changes for object components, capacity devices, and cache tier devices
Object Events	Displays information about vSAN object component state changes

vRealize Log Insight is structured in a hierarchical way allowing you to start an overview level showing events based on a roll up of vSAN components pertinent to the various widgets present on each dashboard. Selecting a rolled up event for a specific component results in displaying the specific, context-sensitive log events for that component.

For example, Figure 6.16 shows the out-of-the-box top Congestion dashboard for vSAN within vRealize Log Insight.

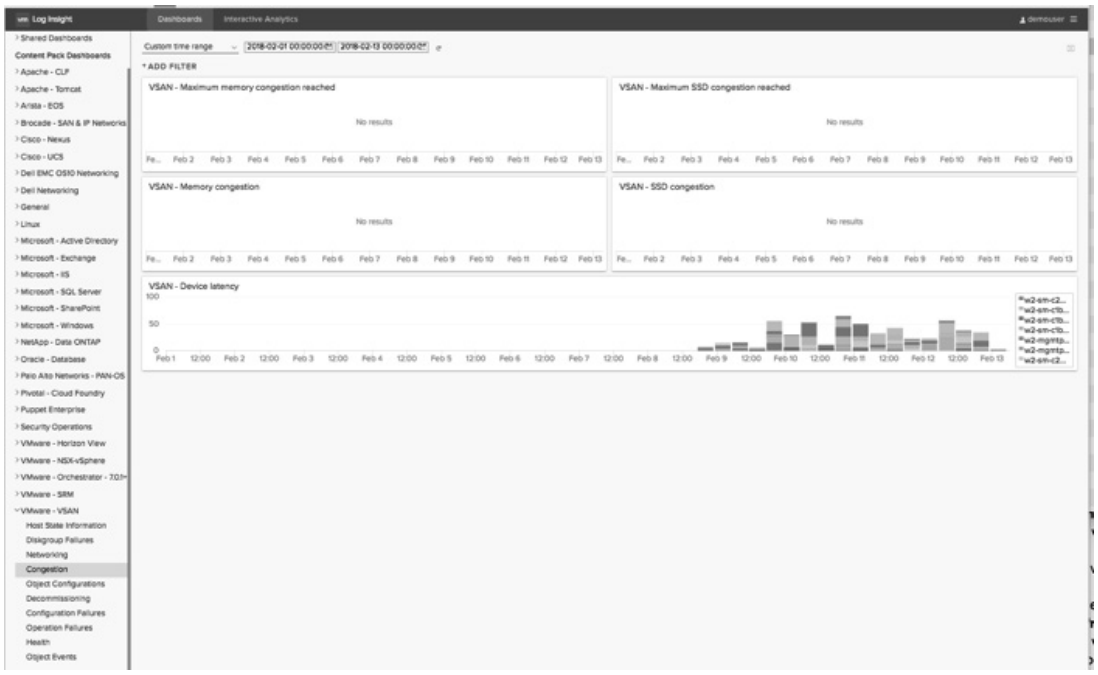

Figure 6.16 vSAN Congestion dashboard in vRealize Log Insight

Selecting a colored portion from one of the bars in the Device latency widget results in the device latency-related log events for that specific object as shown in Figure 6.17.

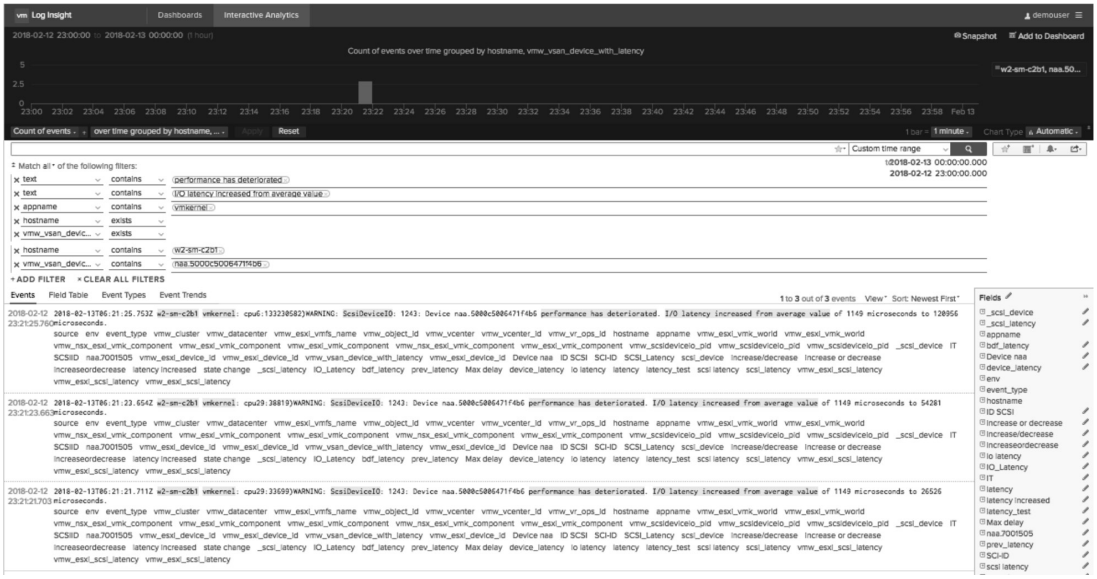

Figure 6.17 Context-sensitive object device latency log events in vRealize Log Insight

In the log entry, we can see for example that the I/O latency for the host selected in the Device latency widget increased from an average value of 1149 microseconds to 120956 microseconds and that it has been steadily increasing over time. While this is shown as a Warning it could be indicative of impending hardware failure or possibly just too many VMs on the host. This is a good example of where further troubleshooting via vRealize Operations might proactively uncover an issue that can be resolved before it becomes service, application, or end-user impacting.

Figure 6.17 above also shows some of vRealize Log Insight's filtering power. In the Filter section you can dynamically extract any field from the data by providing a regular expression. The extracted fields can be used for selection, projection, and aggregation, like how the fields that are extracted at parse time are used. Also, the Fields pane on the right side of the screen shows a list of the data classifications and keywords that you can choose to show or hide. You can also choose to assign a specific value to further filter your results.

Partner ecosystem tools

The VMware Ready for vSAN Program is creating a growing ecosystem of technology partners focused on vSAN integrated file services and data protection tools.

These currently include:

- NetApp ONTAP Select vNAS for VMware vSAN allows clients to present and consume the vSAN datastore capacity via file protocols (NFS/SMB).

- Veeam® Availability Suite™ combines the backup, restore and replication capabilities of Veeam Backup & Replication™ with the advanced monitoring, reporting and capacity planning functionality of Veeam ONE™

Additional technology partners include:

- Commvault for vSAN on Prem providing backup and restore for VMs deployed on vSAN.

- HyTrust DataControl for workload encryption throughout its lifecycle with scalable key management.

  NOTE: For an up-to-date of vSAN ecosystem technology partners, please refer to the VMware Solution Exchange link provided in Table 23 in the "Where you can go for more information" section.

Conclusion

As an integral part of VMware's Software-Defined Data Center, vSAN represents a previously unheard-of level of storage flexibility and agility provided through hyper-converged storage capabilities. But to fully leverage these capabilities in a sustained manner, VMware highly recommends you consider adjusting your operating model, allowing you to optimize the way you operate a software-defined infrastructure. To be truly effective, you can choose to optimize the way you deploy your most valuable resources, your people, in blended teams and ensure they have the right skills to make your vSAN operations successful, and to make them successful as individuals and as a team. You can also choose to optimize your critical storage related operational processes to take advantage of the opportunity hyper-converged storage provides. Lastly, you can leverage the new breed of tools providing intelligent operations to enhance, perhaps even revolutionize, your monitoring and troubleshooting integrated capabilities for vSAN and the associated software-defined data center. Making these changes will produce an operating model that allows you to shift to an intelligent operations mindset. By adopting vSAN technology as well as adapting your operating model in these ways, you will fully leverage your investment in, and realize the benefits provided by software-defined storage.

Where you can go for more information

Table 8.1 References

Description	Reference
Organizing for the Cloud	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutionoverview/vmware-oganizing-for-the-cloud-feb2017.pdf
vSAN Essentials	https://vsan-essentials.gitbooks.io/vsan-6-2/content/
VMware TestDrive – product exploration environment	https://vmtestdrive.com/
vSAN Operations Guide	https://storagehub.vmware.com/#!/vmware-vsan/vsan-operations-guide
vRealize Operations and Log Insight in vSAN Environments	https://storagehub.vmware.com/#!/vmware-vsan/vrealize-operations-and-log-insight-in-vsan-environments
VMware® vSAN Diagnostics and Troubleshooting Reference Manual	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vsan-troubleshooting-reference-manual.pdf
vSAN Health Check Service Information	https://kb.vmware.com/s/article/2114803

Description	Reference
VMware vSAN Top 10 Operational Tips	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vmware-vsan-top-10-operational-tips.pdf
VMware vSAN Quick Monitoring & Troubleshooting Reference Guide	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vmware-virtual-san-quick-monitoring-reference-guide.pdf
VMware vSAN Health Check Plugin Guide	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/products/vsan/vmw-gdl-vsan-health-check.pdf
Using esxcli commands with vSAN	https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-7799D2D7-2513-4372-92EA-4A1FB510E012.html
vSAN API Cookbook for Python	https://storagehub.vmware.com/#!/vmware-vsan/vsan-api-cookbook-for-python
vSAN PowerCLI examples for managing vSAN at scale	https://code.vmware.com/samples?categories=Sample&keywords=vsan&tags=PowerShell
Cost Effective Independent Environments using vSAN	https://storagehub.vmware.com/t/vmware-vsan/cost-effective-independent-environments-using-vsan/
vRealize Operations Management Pack for vSAN (Note: for vSAN 6.6 and above, vRealize Operations Manager includes native vSAN dashboard functionality)	https://marketplace.vmware.com/vsx/solutions/management-pack-for-vsan?ref=search

Description	Reference
vRealize Log Insight Content Pack for vSAN	https://marketplace.vmware.com/vsx/solutions/vmware-vsan?ref=search
VMware Solution Exchange	https://marketplace.vmware.com/vsx/
VMware Storage & Avaliability Technical Documents	https://storagehub.vmware.com/
VMware Hands On Labs	http://labs.hol.vmware.com/
VMware vSAN Support Center	https://www.vmware.com/support/virtual-san.html
VMware Education vSAN Courses	https://mylearn.vmware.com/portals/www/search/results.cfm?ui=www_edu&pID=www&menu=search-results&searchtype=simple&orderBy=relevance&category=catalog&keyword=vSAN&Search=Search
VMware vSAN Product Documentation	https://docs.vmware.com/en/VMware-vSAN/index.html
VMware Storage YouTube channel	https://www.youtube.com/channel/UCOQ1cSf37ags3wnn9XEOC6Q
VMware vSAN Official Product Page	https://www.vmware.com/products/vsan.html
Virtual Blocks Blog	https://blogs.vmware.com/virtualblocks/
VMware Operations Transformation Blog	https://blogs.vmware.com/cloudops
High performance organizations	“Lean Enterprises”, by Jez Humble, Joanne Molesky, and Barry O’Reilly

Index

As VMware's vSphere-native storage solution powering industry-leading Hyper-Converged Infrastructure (HCI) solutions, VMware vSAN® represents a previously unheard-of level of flexibility and agility by providing software-defined, hyper-converged storage capabilities. To fully leverage these capabilities in a sustained manner, certain operational optimizations should be considered.

Operationalizing vSAN brings together the knowledge and guidance for optimizing the on-going operations of the VMware vSAN component of a software-defined data center. It not only addresses tactical optimizations such as tooling for monitoring and troubleshooting but also those of a more strategic nature such as team structure and culture, roles, responsibilities, and skillsets, as well as supporting ITSM process considerations.

Operationalizing vSAN is your roadmap to fully realizing the benefits provided by a software-defined data center running vSAN. You will find insights and recommendations proven in the field for optimizing the way you organize for and operate the environment, unlocking its full potential to provide the flexibility and agility your business stakeholders require.

About the Authors

Kevin Lees is the field Chief Technologist for IT Operations Transformation at VMware, focused on how customers optimize the way they operate VMware-supported environments and solutions. He is responsible for defining, communicating, and evangelizing VMware's IT Operations Transformation vision and strategy as it relates to operational (integrated people, process, and application of VMware technology) best practices. Kevin also serves as an advisor to customer executives globally for their IT operations transformation initiatives.

Paul Wiggett is a Senior Technical Operations Architect working as part of the Transformation Services Practice at VMware. He works as a trusted advisor to some of VMware's largest customers across Europe, the Middle East and Africa assisting and guiding them in their digital transformation strategy and the subsequent transformation of their IT operations teams to a service-oriented and software defined approach.

vmware® PRESS

ISBN-13: 978-0-692-19727-1

Cover design: VMware

$19.95
ISBN 978-0-692-19727-1
51995>

www.vmware.com/go/operationalizing-vsan

9 780692 197271

$19.95