



White Paper

# The Storage Efficiency Guide for NetApp V-Series: Get More From Your Existing Storage

December 2012 | WP-7178

## EXECUTIVE SUMMARY

NetApp® V-Series controllers help you address difficult storage problems such as eliminating storage sprawl in virtual environments, replicating between heterogeneous storage arrays, and reducing storage needed for dev/test while preserving your storage investment. Your existing storage arrays immediately benefit from advanced NetApp storage efficiency technologies including thin provisioning, cloning, replication, deduplication, Snapshot™ copies, and data compression. These technologies increase your effective utilization, reduce capacity requirements by 35% or more, and allow you to manage more storage with fewer resources so they pay for themselves in nine months or less. In addition to dramatic improvements in storage efficiency, V-Series delivers the full suite of NetApp capabilities, including simplified management, unified storage, and integrated data protection.

## TABLE OF CONTENTS

<b>1</b>	<b>INTRODUCTION .....</b>	<b>3</b>
<b>2</b>	<b>HOW THE V-SERIES WORKS .....</b>	<b>4</b>
<b>3</b>	<b>THE V-SERIES STORAGE EFFICIENCY ADVANTAGE .....</b>	<b>5</b>
<b>4</b>	<b>ADDRESSING COMMON STORAGE CHALLENGES .....</b>	<b>8</b>
	<b>ELIMINATING STORAGE SPRAWL IN VIRTUALIZED ENVIRONMENTS .....</b>	<b>8</b>
	<b>REPLICATION BETWEEN HETEROGENEOUS STORAGE ARRAYS .....</b>	<b>9</b>
	<b>NAS AND SAN GATEWAY .....</b>	<b>10</b>
	<b>ACCELERATING THE PERFORMANCE OF EXISTING STORAGE .....</b>	<b>11</b>
	<b>PROVIDING DATA COPIES FOR DEVELOPMENT AND TEST .....</b>	<b>12</b>
<b>5</b>	<b>CONCLUSION.....</b>	<b>13</b>

# 1 INTRODUCTION

IT teams work hard to get their jobs done in the face of tight budgets and growing power, cooling, and space constraints. Rapid data growth—sometimes coupled with mergers and acquisitions—has complicated the situation with heterogeneous storage environments that are incompatible, inflexible, and inefficient. Since starting from scratch is rarely an option, it may seem like nothing can improve your storage situation.

NetApp V-Series open storage controllers are designed to address difficult-to-solve storage problems, improve efficiency, and increase the flexibility of your storage operations. V-Series controllers act as a unified front end for your existing storage arrays, extending the full range of NetApp capabilities, including our industry-leading storage efficiency technologies, to your current storage capacity. This allows you to:

- Achieve 100% or greater storage utilization.
- Reduce capacity requirements by 35% or more\* without sacrificing performance or resiliency.
- Achieve payback in nine months or less.†

For example, a leading American manufacturer deployed V-Series as a front end to its existing EMC and IBM storage arrays. As a result, the company avoided the addition of three full-time-equivalent employees to manage storage, has seen 86% capacity savings, avoided \$300,000 in new capacity spending, and regularly achieves availability of 99.999%.

Using V-Series with your existing storage simplifies your storage environment and eliminates many storage challenges. Deploying V-Series storage controllers can allow you to:

- Take advantage of Snapshot copies, deduplication, thin provisioning, thin replication, compression, virtual copies, Virtual Storage Tiering, and other technologies to improve storage efficiency.
- Manage twice as much data without adding personnel.
- Reduce the storage sprawl associated with virtual server and virtual desktop environments.
- Replicate data between storage arrays from different vendors.
- Consolidate file services onto existing block-based arrays.
- Leverage SAN protocols like iSCSI and FCoE.
- Increase the throughput of your existing storage.
- Accelerate application development by facilitating dev/test.

This paper includes a variety of real-world examples to demonstrate the advantages that V-Series delivers.

---

\* <http://www.netapp.com/us/solutions/infrastructure/virtualization/guarantee.html>


† <http://www.netapp.com/us/products/storage-systems/v-series/pgp.html>

## 2 HOW THE V-SERIES WORKS

Before exploring how the V-Series addresses specific storage challenges, it's important to understand a little bit about the V-Series controllers themselves. NetApp V-Series controllers provide a unified front end that enables you to manage storage arrays from multiple storage vendors—including EMC, HP/3PAR, Hitachi Data Systems, Fujitsu, and IBM—in a consistent way and with less effort. No matter what the native capabilities of your existing storage, the V-Series gives you the full range of advanced NetApp storage capabilities.

A single V-Series controller can be attached to one or more third-party storage arrays. All arrays can be from a single vendor, or multiple arrays from different vendors can be attached to the same V-Series system. You simply connect your V-Series controller and your storage arrays via an existing SAN. NetApp utilizes the native RAID capabilities of each array, but all day-to-day management and provisioning tasks are performed on the V-Series controller.

Figure 1) V-Series controllers virtualize third-party storage to provide all the benefits of NetApp.



If you have a V-Series attached to arrays from EMC, HP, and Hitachi, that's a single point of management—and a single management interface—that you have to deal with. All back-end storage can be accessed and managed as if it's a single storage system, reducing the management overhead and freeing storage administrator time for more value-added tasks. All those storage arrays immediately benefit from a uniform set of NetApp capabilities, including unified storage, integrated data protection, and industry-leading storage efficiency.

### 3 THE V-SERIES STORAGE EFFICIENCY ADVANTAGE

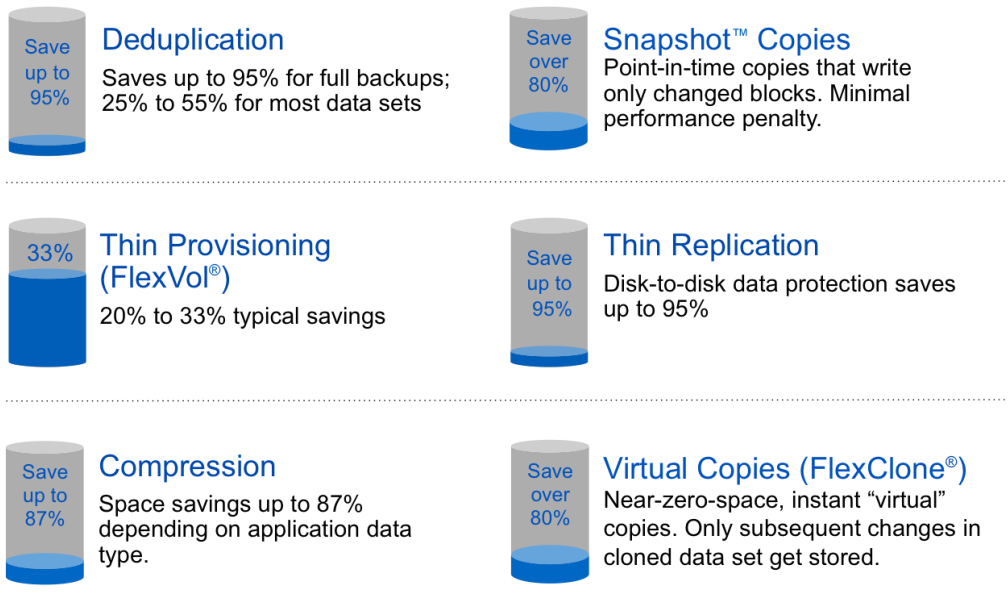
NetApp V-Series controllers let you take advantage of the full range of NetApp storage efficiency technologies. Your third-party storage “inherits” these capabilities when it’s attached to the V-Series controller. This results in much higher storage utilization that frees up storage now and reduces the new capacity you’ll need in the future.

NetApp storage efficiency technologies lower costs, simplify storage operations, and improve data center efficiency. You select which technologies to deploy on your V-Series to fit your specific requirements, and they work together to deliver cumulative savings on the storage you already own.

#### Health and Nutrition Company Uses Several Efficiency Technologies to Achieve Goals

A developer of health and nutrition products reduced storage administration requirements by 87% using V-Series controllers with existing EMC arrays. This company reclaimed over 35TB of capacity with a combination of virtual cloning for its dev/test environment, deduplication, and thin provisioning. Without these technologies the company would have needed about 25% more storage (fully utilized) to meet its storage requirements.

Figure 2) NetApp V-Series storage efficiency technologies and associated savings.



**Snapshot.** A Snapshot copy is a space-efficient, point-in-time copy of a volume or LUN. Although your existing array may include technology that claims to provide similar capabilities, NetApp Snapshot copies offer significant advantages in terms of space efficiency and performance. A NetApp Snapshot copy only consumes additional space as changes are made. The lack of performance impact allows Snapshot copies to be taken more often for a more granular recovery. In practice, NetApp users rely on Snapshot technology as a means of simple, frequent, and efficient backup, and it serves as the foundation of many other NetApp technologies.

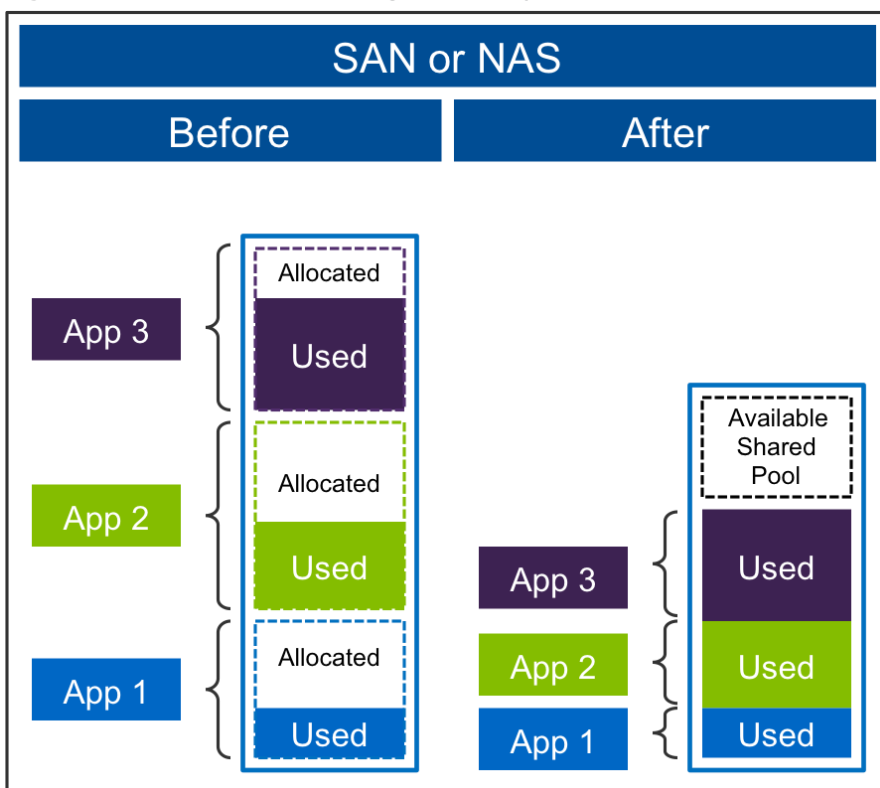
**Deduplication.** Unlike deduplication technologies that work only at the file level, NetApp deduplication works on individual data blocks. NetApp uses a simple “fingerprint” for each block (easy to compute), and

then does a full byte-by-byte comparison when a potential duplicate is found, to verify that two blocks are indeed identical. This approach provides both the efficiency and data integrity necessary to make NetApp deduplication practical for both primary and secondary data. Users routinely employ NetApp deduplication to recover storage capacity consumed by identical data in online backups, virtual environments, home directories, application datasets, and so on. Deduplication recovers up to 90% of the storage space from online backup data, 70% from virtual environments, and 35% from home directories and CIFS shares.

**Thin Provisioning.** When new applications are conventionally provisioned, storage is allocated upfront that cannot be used for anything else. The amount of storage reserved, however, is at best an estimate, and often that estimate is inflated because of the tedious nature of conventional provisioning. The result is that storage you have paid for sits idle and may never be used.

NetApp thin provisioning solves the problem of storage preallocation. Instead of allocating space upfront, storage space is dynamically allocated to each volume or LUN only as data is written. Each application draws from a single, shared storage pool, dramatically reducing the amount of storage you need to purchase. NetApp thin provisioning raises utilization, automates tedious provisioning tasks, and simplifies capacity planning.

Figure 3) NetApp thin provisioning provides “just-in-time” volume and LUN provisioning.



**Thin Replication.** NetApp V-Series controllers provide integrated data protection that allows you to meet your goals more efficiently and with less complexity. NetApp SnapMirror® technology provides efficient, asynchronous, block-level replication to remote locations for disaster recovery. NetApp SnapVault® software provides disk-to-disk backup. Both SnapMirror and SnapVault are based on thin-replication technology that uses both storage and network bandwidth very efficiently. Once a baseline copy is made, only incremental block changes are transferred, not whole files. SnapMirror network compression provides additional bandwidth savings while leaving data on both source and target uncompressed. If

your primary storage is deduplicated and/or compressed, replication consumes even less bandwidth and the resulting secondary volume inherits the space savings.

The NetApp SnapManager suite provides application-aware data protection for popular virtual environments and applications. NetApp SnapProtect™ software combines the advantages of NetApp Snapshot, SnapMirror, and SnapVault with the full cataloging and tape support of more traditional data protection approaches.

**Data Compression.** Until recently, data compression has almost always been done using dedicated hardware that adds to expense and complexity. With the data compression technology built into the Data ONTAP® operating system, NetApp has developed a way to provide transparent inline and postprocessing data compression in software while mitigating the impact on storage system computing resources. You have the flexibility to choose either inline or postprocess compression on your V-Series controllers to meet your specific needs and requirements. Total space savings from applying compression can range up to 87% for some applications.

**Virtual Copies.** There are a variety of situations that require many copies of the same data: application development and test is a common example. In addition, most virtual server and virtual desktop environments include tens or hundreds of virtual machines with the same operating system configuration.

NetApp FlexClone technology allows you to create “virtual copies” of a dataset without requiring a full, time-consuming copy and without doubling the storage needed. A NetApp clone consumes additional disk space only when changes are made. FlexClone lets you make as many copies as you need whenever you need them, while requiring only incremental increases in storage for changed data.

**Virtual Storage Tier.** The NetApp Virtual Storage Tier (VST) provides an integrated, intelligent caching strategy that lets you target deployment of flash acceleration to optimize performance while minimizing cost. VST is dynamic, efficient, and self-managing. Once installed and configured, it requires no ongoing management decisions. Hot data is automatically promoted to the cache with block-level granularity. As access patterns change, algorithms adjust the data that resides in cache automatically to achieve optimal performance.

With NetApp Flash Cache™ intelligent caching, read requests can be satisfied from flash memory inside V-Series controllers 10 times faster than reading data from disk. Flash Cache allows you to satisfy your performance requirements with fewer total disk spindles or with capacity-oriented disk drives rather than more expensive performance-oriented disk drives. Flash Cache works particularly well in conjunction with NetApp deduplication and virtual copies. Because deduplicated blocks are referenced by multiple files, they have a high probability of being reread from cache.

Although Flash Cache provides performance acceleration for your entire V-Series environment, the Flash Accel™ server cache is highly focused. Flash Accel allows data from your V-Series environment to be cached on flash devices inside specific servers, providing maximum throughput and minimum latency for the most performance-sensitive applications.

If you have NetApp disk shelves included as part of your V-Series configuration, you also have the option of including solid-state disks (SSDs) in your NetApp disk shelves and using NetApp Flash Pool™ intelligent caching to accelerate workloads on volumes in the aggregates built on those disk shelves.

**RAID-DP.** While V-Series storage controllers do not require any NetApp disks, you have the option of adding native NetApp disk shelves to your V-Series configuration if you wish. Adding NetApp capacity enables you to use the NetApp dual-parity implementation: RAID-DP® technology. RAID-DP protects you against the possibility of having two disk failures in a single RAID group while incurring almost no performance penalty. Tests show a random write performance delta of only 2% versus the NetApp single-parity RAID implementation. By comparison, another major storage vendor's RAID 6 random write performance decreases by 33% relative to single-parity RAID on the same system. The ability to use RAID-DP instead of RAID 1 (mirroring) for write-intensive workloads reduces the number of disks required and therefore reduces your overall storage costs.

**All Technologies Work Together.** While each of these technologies provides significant benefit by itself and may be deployed individually, they are all designed to work together. You achieve the greatest storage savings by using several capabilities at once.

Using the NetApp AutoSupport™ tool, you'll be able to see how your V-Series controller is performing and how much space you're saving. Our My AutoSupport online tools show you storage efficiency, risk reports with corrective actions, performance and configuration information, and more. [Additional tools](#)‡ can help you benchmark the efficiency of your current IT environment or estimate exactly how much storage you can save.

#### **NetApp Payback Guarantee Program**

NetApp is so sure you'll see significant savings using NetApp storage efficiency technology that we guarantee your V-Series controller will pay for itself in nine months or less—including the software and support required—or we'll make up the difference. This means that by applying storage efficiency technologies described in this white paper, you'll be able to save enough on storage capacity to in effect pay for the cost of the V-Series controller.

Even if you don't sign up for the program, you can track your progress using the My AutoSupport online tool and immediately see how much storage you're saving.

Check out the [Payback Guarantee Web page](#) or [FAQ](#) for more information.

## **4 ADDRESSING COMMON STORAGE CHALLENGES**

The technologies described in the previous sections provide significant and measurable advantages in the real world. Now that you understand the V-Series and the range of capabilities that it brings to your existing storage, let's take a look at how it can help you address your storage challenges.

### **ELIMINATING STORAGE SPRAWL IN VIRTUALIZED ENVIRONMENTS**

Most virtualized environments have dozens or even hundreds of copies of the same operating systems. This high degree of duplication results in significant storage sprawl and makes both virtual server and virtual desktop infrastructure (VDI) environments perfect targets for efficiency improvements.

If your virtualized environment is hosted on a storage array that doesn't support deduplication, adding NetApp V-Series controllers will allow you to reclaim significant amounts of storage that can be reused to meet other needs, reducing your overall storage investment. Typically between 50% and 80% of the storage allocated for virtual machine operating environments can be freed for reuse.

Once a virtualized environment is moved to a V-Series controller, you have several options to maintain the efficiency of the environment. You can continue to use existing provisioning methods combined with deduplication, or you can provision virtual copies of existing virtual machines leveraging the capabilities of NetApp FlexClone.

---

‡ <http://www.netapp.com/us/campaigns/it-efficiency/>

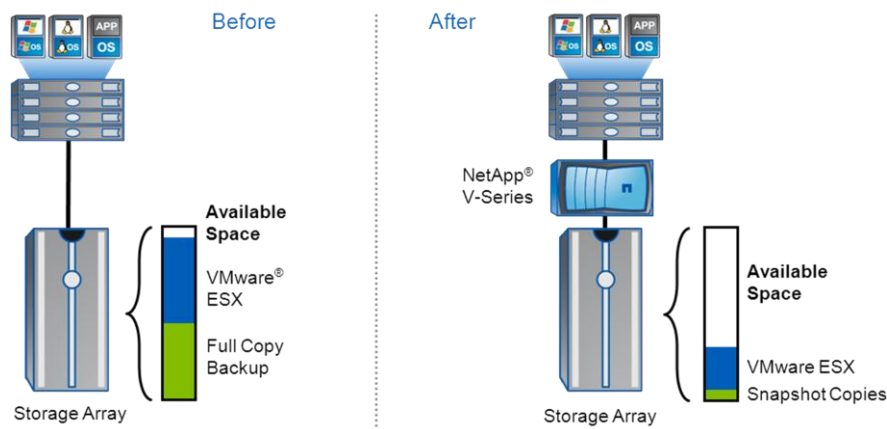


## V-Series and Deduplication Reclaim 86% of VMware Storage

One IT team had a 6TB Hewlett-Packard EVA storage array supporting its VMware environment. Because VM growth and storage consumption had exceeded estimates, only 380GB of free storage space remained in the array. Backups had become inconsistent and difficult to manage.

By using V-Series and NetApp deduplication, the team achieved an 86% rate of deduplication on VMware storage; NetApp Snapshot copies provided additional savings on backups and completely eliminated backup headaches. In combination, over 4TB of space was reclaimed. Free space on the HP EVA storage array increased from 6% to 66%.

Figure 4) VMware storage requirement was reduced by 86% using V-Series with deduplication and Snapshot.

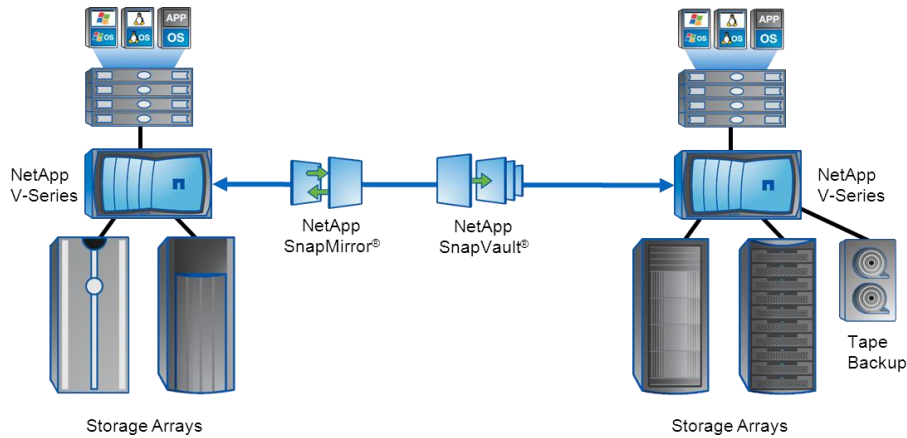


## REPLICATION BETWEEN HETEROGENEOUS STORAGE ARRAYS

Array-based replication solutions for disaster recovery typically require source and target array configurations that are the same. Different array models from the same vendor may not be supported, and if you need to replicate between arrays from two different vendors, all bets are off.

By using V-Series as a front end to your existing storage arrays, you can address these problems and significantly reduce bandwidth requirements. SnapMirror replicates only blocks that have changed, increasing replication efficiency. Deduplicating primary data further reduces the amount of data that has to be replicated, and SnapMirror compression can be enabled to save additional network bandwidth.

Figure 5) V-Series enables replication between heterogeneous storage arrays using either SnapMirror for disaster recovery or SnapVault for disk-based backup.



### Insurance Company Solves Replication and Bandwidth Problems with V-Series

A large insurance company was using LSI storage replication between two IBM storage arrays of different models located 3,000 km (1,800 miles) apart to replicate 200GB of data. Data growth would have required a bandwidth upgrade and the addition of dedicated compression devices.


By replacing the production IBM storage array with a NetApp FAS system and putting a V-Series controller in front of the array at the DR site (so that NetApp SnapMirror could be used), this company reduced the amount of data that needed to be replicated on a daily basis to about 18GB, which was easily accommodated by its existing WAN.

### NAS AND SAN GATEWAY

Sometimes you have capacity available on block-based storage arrays, but you need NAS. One possibility is to configure traditional file servers, but this can result in server sprawl, high latency, and poor reliability. In other situations you may want or need to use a SAN protocol such as iSCSI or FCoE that isn't supported by your existing array.

With the ability to serve both SAN and NAS protocols at the same time, NetApp V-Series is the ideal gateway. A single V-Series controller can take the place of large numbers of file servers, reducing cost and complexity while increasing availability. At the same time it can also deliver block data via FC, FCoE, and iSCSI.

Figure 6) A V-Series controller can act as both a NAS and a SAN gateway, providing file services using your block-based SAN storage or bridging between different SAN protocols.



### University Healthcare Facility Replaces 30 File Servers with Single V-Series

A large university healthcare facility replaced 30 file servers with a highly available NetApp V-Series system front-ending an HDS SAN. This change reduced costs by over \$400K, increased storage utilization by 30%, and saved 750 man-hours per year. At the same time, downtime dropped to zero.

### ACCELERATING THE PERFORMANCE OF EXISTING STORAGE

It may seem counterintuitive that putting a V-Series in front of an existing storage array yields better performance. However, because of the additional level of controller cache—and NetApp technologies such as write anywhere file layout (the WAFL<sup>®</sup> file system) and Virtual Storage Tiering—V-Series can actually drive significantly more I/O operations per second (IOPS) per drive than the same array by itself.

In a [recent study](#)<sup>§</sup>, the Enterprise Strategy Group demonstrated that a V-Series V3270 in front of a third-party array was able to deliver three to six times more IOPS at much lower latency than the array alone on common workloads such as OLTP, Exchange, file services, and Web services. In all cases the datasets were deduplicated, and performance gains were due in part to deduplication-aware caching on the V-Series. Adding Flash Cache to the V-Series further increased performance with gains of five to eight times over the third-party array by itself.

V-Series with Flash Cache installed can be particularly beneficial in VDI environments in which boot storms and login storms can swamp a traditional storage array.

§ <http://media.netapp.com/documents/wp-esg-vseries.pdf>

## Healthcare Provider Increases Throughput More Than 7 Times with V-Series, Flash Cache

A healthcare payment solution provider added several V-Series controllers with Flash Cache and a Fibre Channel over Ethernet (FCoE) fabric as a front end to its EMC and HDS infrastructure. This move not only extended the useful life of the existing storage arrays, it increased maximum throughput from 200MB/sec to 1.5GB/sec, reduced network costs by 30–50%, reduced the amount of data by 50% via deduplication, and saved \$60,000 through deferred purchases of new capacity.

## Telco Accelerates Backup and Restore, Eliminates Tape with V-Series


With exploding data volumes, backup and restore performance may be as big a concern in many data centers as application I/O. A large telco had over 1.5PB of data and was adding 2TB per week. Backups were not completing inside backup windows and restores took days. By putting V-Series controllers in front of existing HDS storage, the IT team achieved 100% backup success, accomplished restores in hours rather than days, and eliminated 80% of its tape requirement.

## PROVIDING DATA COPIES FOR DEVELOPMENT AND TEST

Development and test, quality assurance, data warehousing, and other operations in database environments typically require many copies of production databases. These consume significant amounts of storage—not to mention power, cooling, and data center space. In addition, the time needed to create a full database copy can be prohibitive. These limitations often dictate the way work gets done, as developers and testers struggle to make do with too few copies and datasets that are out of date.

The virtual copy capability supported by NetApp V-Series controllers with NetApp FlexClone lets you correct this situation, reducing the amount of storage needed to a fraction of what would otherwise be required and allowing you to create a database “copy” in minutes.

Figure 7) Using V-Series and FlexClone to streamline the creation of database copies.



## Large Telco Reduces Storage Needed for Dev/Test by Almost 90% with V-Series

A large telecommunications company was struggling with its 4TB CRM database stored on HDS storage. Forty or more copies of the production database were needed at any given time to satisfy various requirements, resulting in the need for 160TB of storage. Making a single copy took in excess of 24 hours. Testing limitations were slowing time to market for application updates. Backup was slow, complex, and expensive and restores were taking more than a day.

The company already had a mirrored copy of the production database on a separate HDS system for DR purposes. Putting a V-Series controller in front of each HDS storage array allowed the company to continue to provide DR while creating the basis for a complete solution to the database copy problem.

Using NetApp FlexClone technology on the mirror copy, the IT team was able to create as many copies as it needed using just a small fraction of the space formerly required. The storage requirement was reduced to just 18TB—an 89% reduction in disk space needed. At the time this corresponded to a disk cost savings of \$15 million and power savings of roughly \$138K, and it reduced the space requirement from 25 cabinets to just 4 cabinets.

Developers can now create a usable clone in less than two minutes, so time to test has gone down while quality has gone up. Snapshot backups are accomplished in two minutes, and restores can be done in two minutes plus the time needed to replay the transaction logs. Backup costs were reduced by over \$2 million per year.

## 5 CONCLUSION

Third-party storage often lacks the advanced storage efficiency features of NetApp storage. Configuring NetApp V-Series open storage controllers as a front end to arrays from EMC, HP, HDS, IBM, and Fujitsu allows you to take advantage of the full suite of NetApp storage efficiency features, including Snapshot, deduplication, thin provisioning, thin replication, data compression, virtual copies, and the Virtual Storage Tier. Using these technologies reduces your storage requirement by at least 35% and savings of 50% or more are common.

But V-Series does more than just improve efficiency; it helps you address your most difficult storage challenges. V-Series can reduce storage and file server sprawl while increasing the performance of your existing storage. It can enable replication between heterogeneous arrays, serve as a gateway so you can use NAS and SAN protocols that your native arrays don't support, and streamline application development so you can deliver new releases faster and with higher quality.

If you're struggling to cope with the complexities of a heterogeneous storage environment, V-Series can help you reduce complexity, decrease administrator time, and cut costs—all without a forklift replacement of your existing storage.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

[Go further, faster®](#)

©2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, AutoSupport, Data ONTAP, Flash Accel, Flash Cache, Flash Pool, FlexClone, FlexVol, RAID-DP, SnapMirror, SnapProtect, Snapshot, SnapVault, and WAFL are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. VMware is a registered trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. WP-7176-1212



[www.netapp.com](http://www.netapp.com)