VMware vSphere® Storage DRS™ Interoperability

TECHNICAL MARKETING DOCUMENTATION
V 1.0/UPDATED MAY 2012

vmware®

Table of Contents

Purpose and Overview

This document presents an overview of best practices for customers considering the implementation of VMware vSphere® Storage DRS™ in combination with advanced storage device features or other VMware® products.

VMware vSphere Storage DRS Introduction

Virtual machine storage provisioning historically has imposed operational challenges. Monitoring datastore capacity and I/O load has proven to be very difficult and as a result is often neglected. During the provisioning process for virtual machines, virtual disk datastores are often randomly selected, leading to hot spots and over- or underutilized datastores. Storage DRS is a new feature introduced in VMware vSphere® 5.0 that provides smart virtual machine placement and load balancing mechanisms based on I/O and space capacity. It will help decrease operational effort associated with the provisioning of virtual machines and monitoring of the storage environment.

Datastore Clusters

vSphere 5.0 introduces a new VMware® vCenter™ object called a datastore cluster—a collection of datastores aggregated to a single unit of consumption from an administrator's perspective. These clusters form the basis of Storage DRS. When one is created, Storage DRS can manage the storage resources in a comparable manner to how it manages compute resources in a cluster. As with a cluster of hosts, a datastore cluster is used to aggregate storage resources, enabling smart and rapid placement of the virtual disk files of a virtual machine and the load balancing of existing workloads. Figure 1 depicts a datastore cluster of 12TB formed by four 3TB datastores.

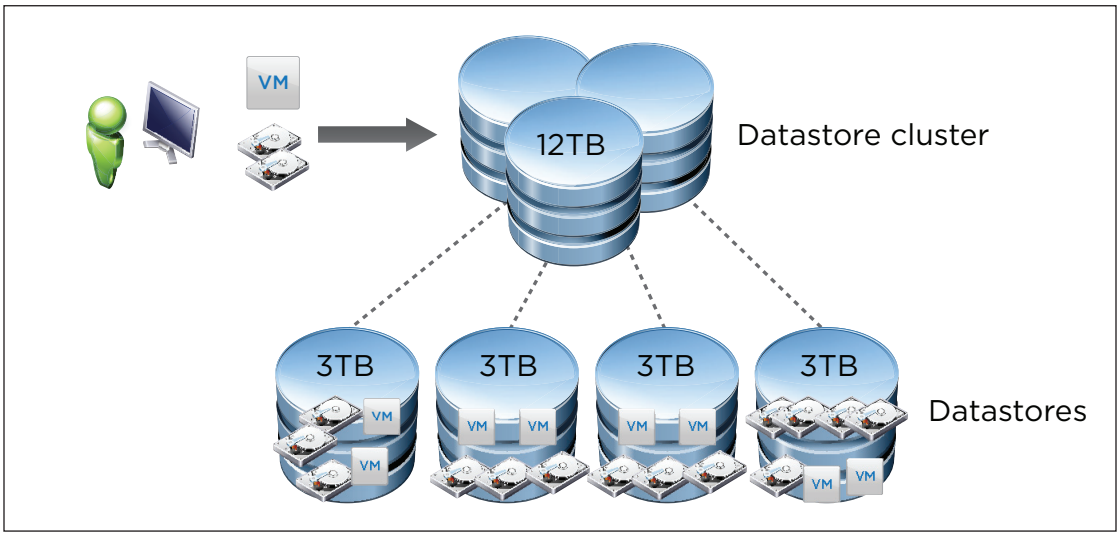

Figure 1. Datastore Cluster

Placement Recommendations

Initial placement in a manual provisioning process has proven to be very complex in most environments. As such, crucial provisioning factors such as current space utilization and I/O load are often ignored. Storage DRS provides initial placement and ongoing balancing recommendations, helping VMware vSphere® administrators make placement decisions based on space and I/O capacity. Storage DRS ensures that initial placement recommendations are made in accordance with space constraints and with respect to the goals of space and I/O load balancing. These goals aim to minimize both the risk of storage I/O bottlenecks and the performance impact on virtual machines.

Storage DRS makes ongoing balancing recommendations when one or more datastores in a datastore cluster exceed the user-configurable space utilization or I/O latency thresholds. These thresholds are typically defined during the configuration of the datastore cluster. Storage DRS uses the datastore utilization–reporting mechanism of VMware vCenter Server™, to make recommendations whenever the configured utilized space threshold is exceeded. I/O load is evaluated by default every 8 hours. When the configured maximum space utilization or the I/O latency threshold (15ms by default) is exceeded, Storage DRS calculates all possible moves to balance the load accordingly while considering the cost and the benefit of the migration.

Affinity Rules and Maintenance Mode

Storage DRS affinity rules enable controlling which virtual disks should or should not be placed on the same datastore within a datastore cluster. By default, a virtual machine's virtual disks are kept together on the same datastore. Storage DRS offers the following three types of affinity rules:

- VMDK Anti-Affinity
  Virtual disks of a virtual machine with multiple virtual disks are placed on different datastores.
- VMDK Affinity (default)
  Virtual disks are kept together on the same datastore.
- VM Anti-affinity
  Two specified virtual machines, including associated disks, are placed on different datastores.

In addition, Storage DRS offers Datastore Maintenance Mode, which automatically evacuates all virtual machines and virtual disk drives from the selected datastore and relocates them to the remaining datastores in the datastore cluster.

Interoperability

Storage DRS leans heavily on VMware vSphere® Storage vMotion® for migrating virtual machines or virtual disks to different datastores. The migrations are recommended or applied when there is an imbalance from either a space utilization or I/O perspective. Depending on the type of storage system used and its capabilities, both the duration and the implication of the migration will vary.

The following table summarizes the various storage system technologies, VMware products, and recommended Storage DRS settings. Each of these will then be described, along with additional details explaining the recommendation and the possible operational impact. To ensure end-to-end support for your environment, validate these recommendations with those of your storage vendor.

FEATURE/PRODUCT	INITIAL PLACEMENT	MIGRATION RECOMMENDATIONS
Array-Based Snapshots	Supported	Manual Mode
Array-Based Deduplication	Supported	Manual Mode
Array-Based Thin Provisioning	Supported	Manual Mode
Array-Based Auto-Tiering	Supported	Manual Mode (only space load balancing)
Array-Based Replication	Supported	Manual Mode
VMware vSphere Raw Disk Mapping	Supported	Fully Automated Mode
VMware vSphere Replication	Not Supported	Not Supported

FEATURE/PRODUCT	INITIAL PLACEMENT	MIGRATION RECOMMENDATIONS
VMware vSphere Snapshots	Supported	Fully Automated Mode
VMware vSphere Thin-Provisioned Disks	Supported	Fully Automated Mode
VMware vSphere Linked Clones	Not Supported	Not Supported
VMware vSphere Storage Metro Cluster	Supported	Manual Mode
VMware vSphere (pre-5.0) Hosts	Not Supported	Not Supported
VMware® vCenter™ Site Recovery Manager™	Not Supported	Not Supported
VMware vCloud® Director™	Not Supported	Not Supported

Table 1. Storage DRS Interoperability

Before we describe each of the features/products separately, the following paragraphs define manual mode and fully automated mode:

In **manual mode**, Storage DRS will make recommendations when the configured thresholds for latency or space utilization have been exceeded. It also will provide recommendations for placement during the provisioning process of a virtual machine or a virtual disk. In other words, when you set Storage DRS to manual mode, you will benefit from it because it will monitor your environment for you and, based on that, will make recommendations on destinations for virtual disk migrations. Manual mode is the default setting when configuring a Storage DRS datastore cluster.

In **fully automated** mode, Storage DRS will make and apply the recommendations automatically when the defined thresholds have been exceeded. By default, Storage DRS is invoked for I/O load balancing every 8 hours or when vCenter has detected that the specified space utilization threshold has been exceeded.

Array-Based Snapshots

Array-based snapshot technology is used in many environments as a mechanism to back up datastores. It can be used safely in conjunction with Storage DRS when the implication of applying a recommendation is fully understood. Storage DRS does not detect array-based snapshots, so they will not be a factor in making or applying recommendations.

VMware recommends configuring Storage DRS in manual mode with I/O metric enabled.

Array-Based Deduplication

Deduplication is currently a more commonly used technique to reduce overall storage costs. When using Storage DRS in conjunction with deduplication, you must determine whether the deduplication process will be as efficient—that is, able to deduplicate the same amount of data—after the migration as it was before the migration. There is always a chance that this is not the case, and this might be a reason not to apply a recommendation to migrate a virtual machine with a large virtual disk. In most scenarios, however, applying all recommendations is advised, to avoid a severe imbalance in latency or an "out of space" condition, which might result in downtime for your environment. Apply these recommendations as close to the start of the deduplication process as possible, or run the deduplication process manually, allowing the storage device to reclaim duplicate blocks.

VMware recommends configuring Storage DRS in manual mode with I/O metric enabled.

Array-Based Thin Provisioning

Array-based thin provisioning poses a minor operational challenge. LUNs presented by "thick-provisioned" storage arrays have their required physical disk space allocated and set aside in advance. When using a thin-provisioned LUN, its logical size as reported by the storage array might be larger than the amount of physical disk space initially allocated to it. Storage DRS by itself does not detect whether the LUN is thin or thick provisioned; it detects only the logical size of the LUN. Provisioning virtual machines on thin-provisioned disks can lead to write operations to previously unallocated blocks of a thin-provisioned LUN, which increases the amount of committed disk space. Because the array has a finite amount of physical storage, it can run out at some point. For this scenario, Storage DRS leverages the VMware vSphere® APIs – Storage Awareness (VASA) thin-provisioning threshold. If the datastore exceeds the thin-provisioning threshold of 75 percent, VASA raises the thin-provisioning alarm. This causes Storage DRS to mark the datastore, which disallows any virtual disk migration to that datastore.

Thin-provisioning alarms are VASA provider specific. Some storage vendors set them on all LUNs of a pool; others set them per LUN or subset of it. Although VMware vSphere® Storage APIs – Array Integration (VAAI) provide a mechanism to reclaim dead space, VMware currently recommends disabling this due to performance issues regarding the reclamation of dead space (VMware knowledge base article 2007427). When applying recommendations, validate that the pool of free disk space is not depleted.

VMware recommends configuring Storage DRS in manual mode with I/O metric enabled.

Array-Based Auto-Tiering

Auto-tiering solutions migrate LUN segments (chunks) to different disk types, based on the use pattern. Hot (frequently accessed) segments typically move to faster disks, whereas cold segments move to slower disks. Depending on the array type and vendor, there are different policies and thresholds for these migrations.

Storage DRS generates migration recommendations that prevent out-of-space situations or resolve imbalances in performance. To prevent hot spots in the datastore cluster and decrease overall latency imbalance, Storage DRS I/O load balancing uses device modeling and virtual machine workload modeling. Device modeling helps Storage DRS analyze the performance characteristics of the devices backing the datastores. Virtual machine workload modeling analyzes workloads of virtual machines running inside the datastore cluster. Both device and workload modeling help Storage DRS assess improvement in I/O latency achieved after a virtual machine migration.

To characterize the datastore, device modeling opens and reads random blocks. Storage DRS does not detect active auto-tiering structures and might open blocks located on disks of similar speed, either slow or fast. The datastore can be distributed throughout multiple tiers of disk types, thereby leading to an incorrect performance characterization. Therefore, the assessment of migrating workload can be incorrect. This can generate a recommendation to move into an underperforming datastore. It also can cause a recommendation to be withheld while there is sufficient performance available in other datastores.

By default, Storage DRS is invoked every 8 hours and requires performance data over more than 16 hours to generate I/O load balancing decisions. Multiple storage vendors offer auto-tiering solutions, each using different time cycles to collect and analyze workload before moving LUN segments. Some auto-tiering solutions move chunks based on real-time workload; other arrays move chunks after collecting performance data for 24 hours. The misalignment of Storage DRS invocation and auto-tiering algorithm cycles makes it unpredictable when LUN segments might be moved, potentially conflicting with Storage DRS recommendations.

By aggregating datastores backed by auto-tiering solutions into a datastore cluster, Storage DRS provides simplified storage management on the datastore cluster level rather than on the datastore level. Initial placement speeds up the process for the user by automating the selection of an individual datastore, leaving the user with the much smaller–scale decision of selecting a datastore cluster. Datastore maintenance provides assistance in storage maintenance tasks by recommending appropriate datastores for evacuating all relevant virtual machines. Load balancing based on space utilization helps prevent out-of-space scenarios and helps balance storage space utilization across datastores inside the datastore cluster.

VMware recommends configuring Storage DRS in manual mode with I/O metric disabled.

Array-Based Replication

In an environment where replication is used, care should be taken when balancing recommendations that are applied. When migrating protected virtual machine files from the source datastore to the destination datastore with Storage vMotion, data must be replicated by the source storage array to the replication partner disk located on the destination storage array. Depending on the method of replication, virtual machine files temporarily can be in an unprotected state. Therefore, migrating protected virtual machines during a scheduled maintenance window is advised.

VMware recommends configuring Storage DRS in manual mode with I/O metric enabled.

VMware vSphere Storage Metro Clustering

VMware vSphere Metro Storage Cluster (vMSC), or stretched cluster as it often is referred to, is an environment that spans multiple sites within a metropolitan area (typically up to 100KM). Storage systems in this type of environment typically enable a seamless failover between sites. These capabilities are provided through the use of synchronous replication.

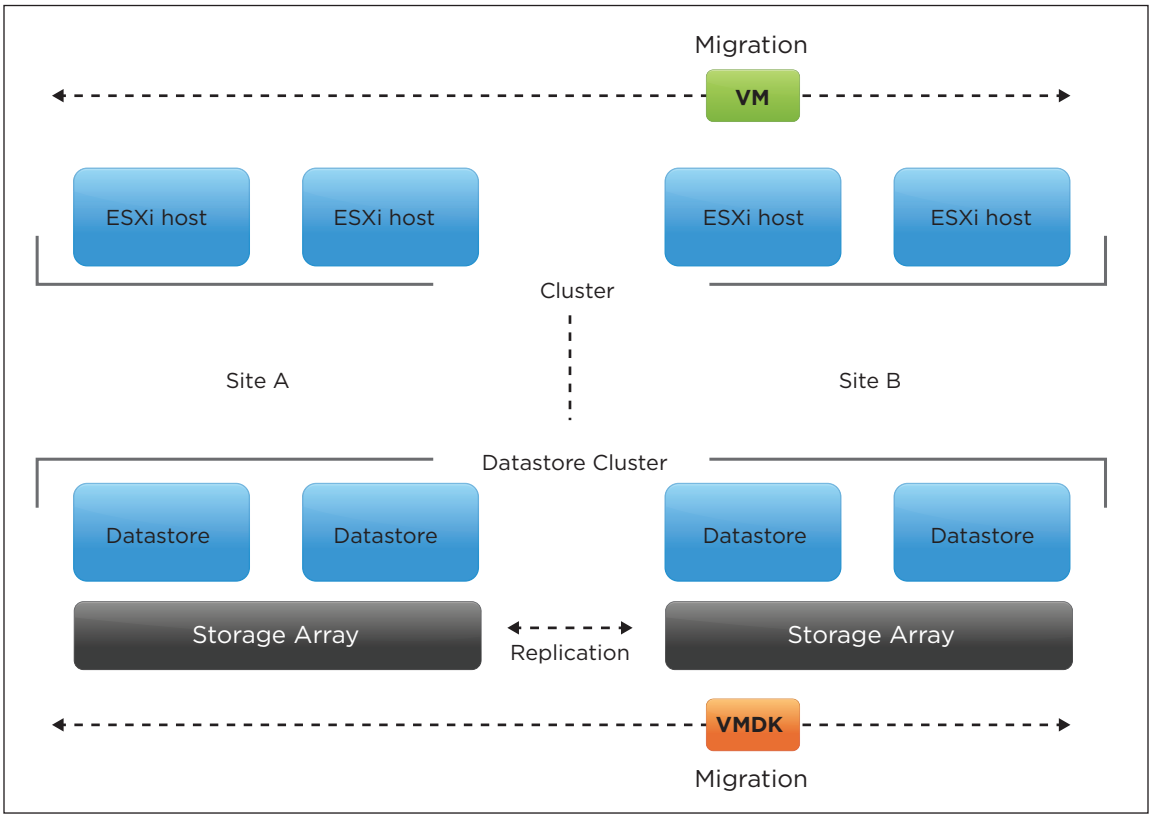

Figure 2. Intersite Virtual Machine Migrations

Care should be taken when aggregating datastores located in multiple sites into one datastore cluster. Storage DRS does not detect site location and can migrate virtual machine files to a datastore that is not local to its VMware® ESXi™ host. This means that data access is through the storage controller in the remote site, through fabric interconnect, possibly increasing latency. If datastores are replicated, a migrating protected virtual machine can temporarily be in an unprotected state, as described in the "Array-Based Replication" section.

VMware recommends configuring Storage DRS in manual mode with I/O metric enabled.

VMware vCenter Site Recovery Manager

Storage DRS does not detect Site Recovery Manager configuration details. If it makes a change for a particular virtual machine, this change might not be in line with what Site Recovery Manager is set up to protect.

Storage DRS does not detect Site Recovery Manager datastore groups and does not take them into account when migrating a virtual machine into another datastore. Because Site Recovery Manager does not detect migrations of virtual machines between datastores, it is not able to update the protection groups automatically. This can lead to a situation in which the virtual machine is left unprotected and potentially violates the recovery plan. Similar situations occur when datastores in the datastore cluster are not used by Site Recovery Manager. Recovery plans might again be violated if Storage DRS moves a virtual machine to such a datastore.

VMware cannot ensure that Site Recovery Manager configurations will not be violated after applying a migration recommendation. Therefore, neither vSphere Storage vMotion nor Storage DRS is supported in combination with virtual machines protected by Site Recovery Manager.

VMware vSphere Raw Disk Mapping

vSphere 5.0 fully supports the use of Raw Disk Mapping (RDM) in combination with Storage DRS. The target LUN that is dedicated to an RDM disk will not be managed by Storage DRS, but it will manage the pointer mapping file for an RDM disk. Pointer files are managed in the same way as the other disk files of a virtual machine. Storage DRS is able to place a virtual machine with an RDM pointer mapping file, as well as relocating the virtual machine. Cloning virtual machines and adding disks with pointer files are also supported by Storage DRS. If a datastore is placed into maintenance mode, Storage DRS will migrate the pointer file to another datastore in the datastore cluster. If the virtual machine is configured with a VMDK anti-affinity rule, it will be unlikely for Storage DRS to move the RDM pointer file for space balancing because the pointer file is very small. In addition, I/O load balancing is effectively disabled for RDM because the pointer files themselves are not involved in I/O, and the I/O to the target LUNs cannot be redirected by Storage DRS.

VMware supports vSphere Raw Disk Mapping in Storage DRS.

VMware vSphere Replication

VMware vCenter Site Recovery Manager 5.0 introduces vSphere Replication, a feature that enables vSphere 5.0 hosts to perform replication of powered-on virtual machines over the network to another vSphere 5.0 host, without the requirement for storage array–based replication.

vSphere Replication tracks changes and replicates the changed blocks to a shadow copy of the virtual machine residing at the recovery site, in accordance with a recovery point objective set as a property of the virtual machine itself.

When applying the default affinity rule inside a datastore cluster, all virtual machine files detected by Storage DRS will be migrated to another datastore if a migration recommendation is applied. Persistent state files (".psf") introduced by vSphere Replication are not detected by Storage DRS and are not migrated, but rather deleted, when a virtual machine is moved. If vSphere Replication does not detect a ".psf" file, it treats this as a power-off operation of the virtual machine and a power-on without the ".psf." This requires it to start a "full sync" operation, which reads the entire disk and validates the integrity of the data. This is done on both the source and destination datastore. Although primary and replica data are in sync, this process can take some time, potentially delaying replication. This might lead to missing recovery point objective windows and a scenario where the virtual machine, when recovered, might be more out of date than had been planned.

Storage DRS does not recognize the shadow copy data and the persistent state file. Therefore, it cannot migrate this data out of the datastore if the datastore is placed in maintenance mode or exceeds the space utilization threshold.

VMware does not support Storage DRS in combination with vSphere Replication.

VMware vSphere Snapshots

Storage DRS supports virtual machine snapshots. By default, it collocates them with the virtual machine disk file to prevent fragmentation of the virtual machine. Also by default, Storage DRS applies a VMDK affinity rule to each new virtual machine. If it migrates the virtual machine to another datastore, all the files, including the snapshot files, move with it. If the virtual machine is configured with an inter-VMDK affinity setting, the snapshot is placed in the directory of its related disk and is moved to the same destination datastore as when migrated by a Storage vMotion operation.

VMware supports the use of vSphere snapshots in combination with Storage DRS.

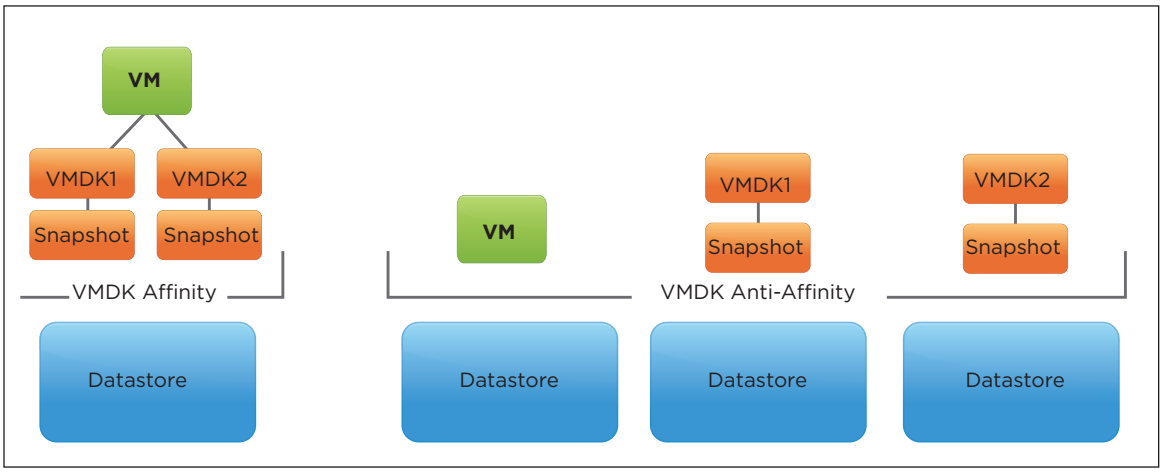

Figure 3. Snapshot Placement with VMDK Affinity and VMDK Anti-Affinity Rules

VMware vSphere Thin-Provisioned Disks

Storage DRS supports thin-provisioned virtual machine disk files. It takes consumed space, as opposed to provisioned space, into account during load balancing operations. The monitoring of growth rates helps Storage DRS detect the impact of moving a thin-provisioned virtual machine disk file. When a given file placement in combination with the estimated growth rate will result in exceeding the threshold in the following 30 hours, Storage DRS will refrain from recommending this move.

VMware supports the use of vSphere thin-provisioned disks in combination with Storage DRS.

VMware vSphere Linked Clones

Storage DRS does not support vSphere linked clones. Because the Storage DRS model does not detect the sharing relationship between two linked-clone virtual machines, a potential danger exists in the loss of space efficiency when moving linked clones. To protect the integrity of linked-clone virtual machines, Storage DRS rejects placement operations that specifically create linked clones.

VMware does not support the use of vSphere linked clones in combination with Storage DRS.

VMware vCloud Director

VMware does not support connecting vCenter Server 5.0 clusters used with VMware vCloud® Director™ to Storage DRS datastore clusters. If vCloud Director makes an attempt to place a VMware vSphere® vApp™ on a datastore belonging to a datastore cluster, the operation will fail.

VMware does not support the use of Storage DRS in combination with vCloud Director.

VMware vSphere (pre-5.0) Hosts

Storage DRS does not support vSphere (pre-5.0) hosts connected to a datastore in a datastore cluster. If older versions of vSphere or VMware® ESX®/VMware ESXi hosts are attached to a VMware vSphere® VMFS3 datastore in a datastore cluster, Storage DRS load balancing will be disabled. No migration recommendations will be made until all hosts in the datastore cluster are vSphere 5.0 hosts.

VMware does not support vSphere (pre-5.0) hosts connected to Storage DRS–enabled datastores.

Conclusion

VMware vSphere 5.0 Storage DRS can help decrease the amount of operational effort associated with monitoring your storage environment and provisioning virtual machines and disks. Various advanced array features and VMware products are used in many such environments. It is important to realize the impact that virtual machine or virtual disk migration can have in these scenarios.

About the Authors

Frank Denneman is senior architect in the Technical Marketing group at VMware. His focus is on resource management and downtime avoidance. He specializes in vSphere DRS, Storage DRS, vMotion, and Storage vMotion. He is the co-author of multiple books including best sellers *vSphere 5.0 Clustering Technical Deepdive* and *VMware vSphere 4.1 HA and DRS Technical Deepdive*. Frank is the owner and main author of leading virtualization and VMware blog frankdenneman.com.

- Follow Frank's blogs at http://www.frankdenneman.nl and http://blogs.vmware.com/vSphere.
- Follow Frank on Twitter: @frankdenneman.

Duncan Epping is principal architect in the Technical Marketing group at VMware. His focus is on cloud infrastructure architecture and integration. Previously he worked at Oracle and several consultancy companies, where he had more than 10 years of experience designing and developing infrastructures and deployment best practices. He was among the first VMware Certified Design Experts (VCDX 007). He is the co-author of multiple books, including best sellers *vSphere 5.0 Clustering Technical Deepdive* and *VMware vSphere 4.1 HA and DRS Technical Deepdive*. Duncan is the owner and main author of leading virtualization and VMware blog Yellow-bricks.com.

- Follow Duncan's blogs at http://www.yellow-bricks.com and http://blogs.vmware.com/vSphere.
- Follow Duncan on Twitter: @DuncanYB.

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