


VMware vSphere® Flash Read Cache® 1.0

VMware vSphere Flash Read Cache



Host Configuration

- Q1. What are the minimum and maximum numbers of hosts required for VMware vSphere® Flash Read Cache®?**
- A1. A minimum of 1 and a maximum of 32 hosts per cluster.
- Q2. What is vSphere Flash Read Cache VFFS?**
- A2. VFFS is a logical container used by the VMware vSphere platform to group and manage local Flash devices into a single resource.
- Q3. Is there a way to configure the VFFS for multiple VMware ESXi™ hosts simultaneously?**
- A3. Yes. The VFFS can be configured in batch mode from the VMware vSphere Web Client.
- Q4. Is it possible to configure multiple VFFS logical containers per ESXi hosts?**
- A4. No. You can have only one VFFS logical container per ESXi host.
- Q5. Is it possible to utilize the VFFS for ESXi host memory swap?**
- A5. Yes. vSphere Flash Read Cache offers legacy support for the swap-to-SSD feature introduced in vSphere 5.0.
- Q6. To use vSphere Flash Read Cache, which version of vSphere is required?**
- A6. vSphere 5.5 or later.
- Q7. Will vSphere Flash Read Cache be backward compatible with vSphere 5.1?**
- A7. No. vSphere Flash Read Cache is compatible with only vSphere 5.5 and later.

Disk Configuration

Q8. What types of storage device interface are supported for vSphere Flash Read Cache?

A8. SATA, SAS and PCI Express.

Q9. What is the maximum number of Flash devices that can be used per ESXi host for vSphere Flash Read Cache?

A9. A maximum of eight Flash-based devices per ESXi host.

Q10. What is the maximum Virtual Flash Host Swap Cache supported by ESXi hosts?

A10. The maximum supported is 4TB per ESXi host.

Q11. What is the maximum virtual machine Virtual Flash Read Cache working set supported per virtual machine disk format (VMDK)?

A11. The default Virtual Flash Read Cache configurable maximum size is 200GB, but a 400GB total maximum is supported.

Q12. How can the total supported Virtual Flash Read Cache be configured?

A12. To use a 400GB cache working set, the host's advanced settings values require the modification of VFLASH.MaxCacheFileSizeMB to 409600 and VFLASH.MaxDiskFileSizeGB to 16384.

Q13. What are the block-size settings supported by Virtual Flash Read Cache?

A13. The minimum supported block size is 1KB; the maximum supported block size is 1,024KB.

Q14. Is it possible to configure Virtual Flash Read Cache at the virtual machine level?

A14. No. The Virtual Flash Read Cache is granularly allocated on a per-VMDK basis.

Q15. Can vSphere Flash Read Cache be used with remote Flash devices?

A15. No. vSphere Flash Read Cache works with only local Flash devices.

Q16. Does vSphere Flash Read Cache provide write-back and write-through cache modes?

A16. No. The 1.0 version of vSphere Flash Read Cache supports only write-through mode (read only).

Q17. Which types of vSphere datastores are supported for virtual machines using Virtual Flash Read Cache?

A17. VMware vSphere VMFS; Network File System (NFS); Raw-device mapping (RDM).

Q18. Does vSphere Flash Read Cache support the use of physical RDMs?

A18. No. vSphere Flash Read Cache supports the use of only virtual RDMs.

Virtual Machine Configuration**Q19. What are the virtual machine requirements for using Virtual Flash Read Cache?**

A19. The virtual machine hardware version must be upgraded to version 10 (VMX-10) and the vSphere Web Client for management.

Q20. Can the vSphere C# client be used to manage and configure Virtual Flash Read Cache?

A20. No. To configure and manage any of the features and components of Virtual Flash Read Cache, the vSphere Web Client must be used.

Q21. To use Virtual Flash Read Cache, what virtual machine agents are required?

A21. None. Virtual Flash Read Cache is an agentless solution and is completely transparent to virtual machines and guest operating system (OS).

Q22. Do adding and removing Virtual Flash Read Cache settings from a virtual machine require a shutdown?

A22. No. The Virtual Flash Read Cache working set can be added, removed and modified dynamically without having to shut down virtual machines.

Q23. What happens to a virtual machine's cache contents when the settings are modified?

A23. Increasing, reducing or modifying the virtual machine's Virtual Flash Read Cache settings will result in the cache contents' being discarded.

Q24. Can Virtual Flash Read Cache settings be configured for multiple virtual machines simultaneously from the vSphere Web Client?

A24. No, not currently. Configuring multiple virtual machines simultaneously can be achieved via one of our scripting solutions, such as VMware vSphere PowerCLI™.

Q25. Does Virtual Flash Read Cache enhance the performance of all virtual machine applications after it has been configured?

A25. No. vSphere Flash Read Cache is designed to enhance the performance of applications that have I/O patterns based on read-intensive operations.

Q26. Can VMware vSphere Fault Tolerance (vSphere FT) enabled virtual machines use Virtual Flash Read Cache?

A26. No. vSphere FT enabled virtual machines are not supported.

Virtual Flash Resource Management**Q27. What is the maximum number of Flash devices supported per VFFS?**

A27. A maximum of eight Flash devices is supported.

Q28. What is the maximum capacity supported for the VFFS per host?

A28. A maximum capacity of 32TB per VFFS.

Q29. What is the maximum capacity supported per Flash device?

A29. A maximum of 4TB is supported per Flash device.

Q30. How can the resource utilization of the VFFS pool be monitored and tracked?

A30. The VFFS utilization metrics are displayed on a per-host basis under the Virtual Flash settings and at the cluster level as part of the cluster resources under the summary tab.

Q31. Can the VFFS be used to store virtual machines?

A31. No. The VFFS is a caching layer—not a storage device that can be used to store virtual machines. It is not visible in the UI.

Q32. Can the VFFS or a Flash device be shared between VSAN and vSphere Flash Read Cache?

A32. No. Flash devices cannot be shared between VSAN and vSphere Flash Read Cache. Each requires its own exclusive and dedicated Flash device.

Q33. Can the VFFS be used to provide resources simultaneously to ESXi hosts as well as to virtual machines?

A33. Yes. After the VFFS has been created, the total available capacity can be used and consumed by both ESXi hosts as Virtual Flash Host Swap Cache and virtual machines as Virtual Flash Read Cache.

Q34. Is it possible to manually assign dedicated Flash devices to individual VMDKs?

A34. No. The VFFS combines all Flash devices into one entity, and it is not possible to pick and choose individual Flash devices.

Q35. Does the VFFS leverage shares to manage resource contention?

A35. No. The allocation of Virtual Flash Read Cache is treated as a reservation, so there aren't any issues with resource contention.

Q36. Does the VFFS provide support for expandable reservation?

A36. No. There is no support for expandable reservations.

Q37. Does the VFFS support thin provisioning?

A37. No. There is no support for thin provisioning or oversubscription.

Interoperability

Q38. Is vSphere Flash Read Cache compatible with VMware vSphere vMotion*?

A38. Yes. The vSphere vMotion workflow has been modified to include two new migration settings for cache contents:

- Always migrate the cache contents (copy)
- Do not migrate the cache contents (drop)

Q39. The VFFS is a nonshared local resource container. How can vSphere vMotion work?

A39. VFFS is compatible with vSphere vMotion, svMotion and XvMotion. vSphere vMotion migrations are automatically converted into an XvMotion operation in this case.

Q40. What is the impact of migrating the virtual machine cache contents?

A40. Depending on the size of the cache contents, the migration operation can take longer than usual to complete, but the virtual machine maintains the cache contents.

Q41. What happens if the cache contents are discarded as part of the migration operation?

A41. The migration operation will complete much faster, but the data must be recached on the target host after the migration has been completed.

Q42. How does VMware vSphere Distributed Resource Scheduler™ (vSphere DRS) handle migrations of virtual machines with Virtual Flash Read Cache?

A42. vSphere DRS doesn't automatically migrate virtual machines with Virtual Flash Read Cache as part of its load balancing operations, except when a host is placed in maintenance mode or during major overutilization contention events.

Q43. Can the VFFS be managed at a cluster level?

A43. No. The VFFS is managed at the host level only. During virtual machine power-on, vSphere DRS makes initial placement based on the available VFFS on each host at the cluster level.

Q44. Are the Virtual Flash Read Cache contents replicated to other hosts with a VFFS?

A44. No. There is no replication on cache contents across hosts in a cluster.

Q45. How does VMware vSphere High Availability (vSphere HA) handle a virtual machine failure event?

A45. If a virtual machine crashes, vSphere HA will restart the virtual machine on the same hosts, and cache contents will be rebuilt.

Q46. If a virtual machine I/O pattern changes, will the Virtual Flash Read Cache be dynamically updated?

A46. No. In version 1.0, the allocation of Virtual Flash Read Cache is treated as a reservation, which also represents the limit. It is not automatically increased or decreased by the system.

Q47. What functional operations will preserve the cache contents of a virtual machine?

A47. Operations such as snapshot, cloning, migrations, and fast suspend/resume will preserve cache contents.

Q48. What functional operations will discard the cache contents of a virtual machine?

A48. Operations such as suspend, resize, change, delete, modify, virtual machine restarts, host restart, and restore snapshot will discard cache contents.

Q49. When using vSphere Flash Read Cache in a vSphere DRS or vSphere HA enabled cluster, is it required that all hosts have VFFS?

A49. No. vSphere Flash Read Cache does not have to be enabled for every host in a cluster, but it is recommended.

Q50. Is it possible to migrate Virtual Flash Read Cache enabled virtual machines onto hosts without a VFFS?

A50. No. The VMware® vCenter Server™ compatibility resource and compatibility checkpoints will prevent the migration from occurring.

Q51. Because the Virtual Flash Read Cache settings for virtual machines are based on reservations, when is the working set capacity actually consumed?

A51. The VFFS capacity is consumed by virtual machines only when they are powered on and when migrated to another host. When the virtual machines are powered off, the allocated capacity is put back into the VFFS pool.

Q52. Do capacity allocation and consumption for Virtual Flash Host Swap Cache work the same way as for virtual machines?

A52. No. The Virtual Flash Host Swap Cache capacity is consumed immediately. It is never given back to the VFFS regardless of the power state of the host.

Q53. Does vSphere Flash Read Cache provide some capability for selecting what data to cache?

A53. No. vSphere Flash Read Cache does not use any application or guest-level intelligence to selectively choose what data must be cached. It caches all data, based on I/O heuristics maintained by vSphere Flash Read Cache.

Device Failures and Failure Scenarios**Q54. What is the impact to a virtual machine when the Flash device storing its cache contents fails?**

A54. The virtual machine will suffer a performance hit because its cache contents will not be available and all I/O will go directly to the VMDKs.

Q55. Will the failure of a Flash device storing cache contents cause virtual machine outages?

A55. No. Virtual machines will continue to operate without their cache contents, but performance will be impacted.

Q56. What happens when the VFFS capacity is completely utilized?

A56. It will not be possible to assign or reserve any more Virtual Flash Read Cache resources to new virtual machines.

Q57. What happens to a vSphere Flash Read Cache enabled virtual machine that is failed over to another host via vSphere HA when the VFFS on the target host is full?

A57. The virtual machine will fail to power on, due to the lack of spare capacity.

Q58. Will disabling vSphere HA admission control allow virtual machines to power on if there is a failover event for which the receiving hosts cannot satisfy the vSphere Flash Read Cache reservations?

A58. No. The virtual machines will fail to power on. The resources provided by vSphere Flash Read Cache are not taken into account by vSphere HA admission control operations.

Q59. What happens when the Flash device's cache capacity is full?

A59. There is a threshold mechanism that is set to 90 percent of utilization. When the cache fill percentage reaches threshold, the eviction process is triggered.

Q60. What happens with the cache contents when the eviction process is triggered?

A60. Cache contents will start evicting pages, based on the usage history.

Q61. Is there any intelligence applied to the data eviction process?

A61. vSphere Flash Read Cache adapts data usage information obtained by detecting cache hits to previously cached data. It then determines which data should be retained or evicted, based on least impact to the virtual machines and to application performance.

Q62. Can vCenter Server monitor the use of vSphere Flash Read Cache?

A62. Yes. There are three new performance statistics counters included that monitor the use of vSphere Flash Read Cache enabled virtual machines.

Q63. What are the names of the vSphere Flash Read Cache related performance statistics counters, and what types of statistics do they provide?

A63. The following are the performance statistics counters:

- vFlashCachelops – Monitors the average number of I/Os per second to the virtual disk cache.
- vFlashCacheLatency – Monitors the average number of microseconds to complete I/O to the virtual disk cache.
- vFlashCacheThroughput – Monitors the average of cache currently being controlled by the VFC module.

Q64. Does vSphere Flash Read Cache provide any CLI-based monitoring solutions?

A64. Yes. The cache performance statistics can be retrieved in greater detail by using the new vSphere Flash Read Cache ESXCLI namespace.

Q65. What is the performance statistics syntax for the vSphere Flash Read Cache ESXCLI namespace?

A65. `esxcli storage vflash cache stats get -m <module> -c <cache file name>`.

Q66. Where are the vSphere Flash Read Cache files located?

A66. The vSphere Flash Read Cache files are located on the ESXi host under the `/vmfs/volumes/vffs/vflash`.

Hardware**Q67. What Flash devices are recommended for vSphere Flash Read Cache?**

A67. Currently the VMware Compatibility Guide has a list of officially supported devices from Dell, Intel and Samsung. Visit the VMware compatibility site for the latest information on supported devices.

<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=vfrc>

About the Author

Rawlinson is a senior architect working for the Cloud Infrastructure Technical Marketing organization at VMware. His focus is on software-defined storage technologies and OpenStack framework integration with VMware products and solutions. Previously he was an architect in the VMware Cloud Infrastructure and Management Professional Services organization, focused on vSphere and cloud enterprise architectures for VMware Fortune 100 and 500 customers. Rawlinson is among the first VMware Certified Design Experts and the author of multiple books based on VMware and other technologies.

- Follow Rawlinson's blogs:
 - <http://blogs.vmware.com/vsphere/storage>
 - <http://www.punchingclouds.com>
- Follow Rawlinson on Twitter:
 - [@PunchingClouds](https://twitter.com/PunchingClouds)

