Run and Reporting Rules for VMmark Version 3.x

Updated August 22, 2022

1 Introduction

The VMmark® Run and Reporting Rules define how to correctly measure and report performance and (optionally) power using the VMmark 3.x benchmark. This document outlines the requirements for producing a publication-quality VMmark result. Only results that meet these requirements can be published using the VMmark metric.

VMware retains the right to update and re-issue new versions of VMmark, including updates to these rules. See Section 5 of these rules for further detail.

1.1 Definitions

In a virtualized environment, the definitions of commonly used terms can have multiple or different meanings. To avoid ambiguity, this section attempts to define terms that are used throughout this document:

- **Server:** A server is a system that is capable of supporting a single native operating system or hypervisor. A server consists of one or more enclosures that contain hardware components such as processors, memory, network adapters, and storage adapters, as well as the mechanism that provides power for these components. In the case of a blade server, the blade enclosure is defined as part of the server.
- **SUT:** The SUT, or System Under Test, is defined as the servers, the virtualization layer, and performance-critical components that execute the defined workloads and operations of VMmark. In addition to the servers, these components include external storage hardware and all hardware necessary to connect the servers to the storage subsystem, such as Fibre Channel switches (in the case of SAN storage) or network switches (in the case of NAS storage). The data-center management server, client hardware used to initiate and monitor the workflow, and network switches not described above are not considered part of the SUT.
- **Client:** Clients are one or several virtual machines that are used to initiate benchmark transactions and record their completion. A Client simulates the work requests that would normally come from end users. Clients are not part of the SUT.
- **Power monitoring systems:** These are the power meters and system(s) running the applications that control the collection and recording of power information for the benchmark. Power monitoring systems are not part of the SUT.
- **Fully supported:** For a product to be considered “fully supported” for purposes of VMmark results, the product’s vendor or the vendor’s designated alternate entity must be available to assist users with any problems they experience with the product. In cases where the mechanism to obtain assistance differs from a conventional enterprise support arrangement, contact VMware to see if it meets the VMmark support requirement.
- **Publicly Available:** To be considered “Publicly Available” in the context of VMmark, a product must be shipping and the vendor or its designated support channel must fully support

it as it used in the VMmark submission. If a pre-release component is used, the benchmarker must have verified that the vendor is committed to making the component available within the 90-day timeframe. The publication of a VMmark result with a future availability date is considered a public commitment that any potential customer for that component must be able to confirm with the vendor or supplier.

2 VMmark Reporting Rules

The VMmark Reporting Rules section of this document describes the requirements and options for publishing VMmark results. It also describes a review process that VMware strongly encourages benchmarkers to take advantage of prior to publication of a VMmark result. This section also includes a set of fair-use rules for use of VMmark results in competitive comparisons and for use in academic and research papers.

VMmark results can be published in three ways:

- They can be published on the VMmark results page.
- They can be independently published.
- They can be published under the academic and research rules.

This section describes the rules for each of these three publication types.

2.1 Publication on the VMmark Results Page

VMware encourages all VMmark benchmarkers to formally submit a full disclosure report of their VMmark results and supporting documentation to the VMmark Review Panel for review and publication on the VMmark results page. The VMmark Review Panel consists of members from partner companies who regularly publish VMmark results and/or have detailed expertise with the VMmark benchmark. It is intended to maintain impartiality of published VMmark results. The Review Panel members have signed agreements to keep the details of VMmark submittals confidential until they are published. It is recommended that submitters of benchmark results review the *Review Panel Guidelines and Operating Procedures*. A roster of member companies on the Review Panel is available on the VMmark website.

To be published on the VMmark results page, results must first be reviewed by the review panel and accepted by VMware. Alternatively, or in addition, benchmark users can publish independently, as described in Section 2.2.

Details of the VMmark review and publication process:

a) The submitter will deposit the appropriate submission information (as defined in Section 2.3 of this document) to the official VMmark Review Panel ftp site. For the URL and access information send an inquiry email to benchmark@vmware.com.

b) An automatic process runs every other Tuesday (“Day 0”) at 9:00 am Pacific time (regardless of any holidays). This process moves all results posted in the submit FTP site to a location visible to all of the Review Panel members.

c) The submitter must ensure that a full, compliant, and complete submittal is in place before 9:00 am Pacific time on Day 0 of the review cycle.
Note: Preliminary compliance is indicated in the score.txt file; the lack of such indication will automatically cause a result to not be considered for review.

d) The submitter may submit a new submittal as early as 13 days before the start of the intended review cycle. Submitted information will not be visible outside of VMware until Day 0 of the next review cycle.

e) The testing environment should remain unchanged and accessible until the review process has been formally completed.

f) Once a result has been deposited on the submit FTP site, a submitter may withdraw that result prior to the opening of the review cycle window by requesting via email sent to benchmark@vmware.com at least two business days (which would not include U.S. holidays) prior to Day 0 of the next review cycle that VMware remove the result.

g) A submitter may supersede a previously submitted result with another for a particular review cycle by first withdrawing the original submittal as described above and then submitting the replacement result before 9:00 am Pacific time on Day 0 of the review cycle.

h) The review cycle start dates will be every other Tuesday except in the infrequent event of a schedule change by VMware. Any changes will be announced on the VMmark Communities page (https://communities.vmware.com/community/vmtn/performance/vmmark). A cycle ends at the exact same time as the next cycle begins, 9:00 am Pacific Time.

i) On Day 0 of the review cycle, all results submitted will be made available to all Review Panel members no later than 11:00 am Pacific time. Notification will be sent to the Review Panel members via email which will provide an FTP site where the results can be found.

j) Review Panel members have seven calendar days to flag any issues via email and are encouraged to flag issues as early as detected. On approximately Day 7 of the review cycle there will be a Review Panel phone conference to discuss any flagged issues or concerns among members and forward them to the submitter as needed. The exact date of the conference call may be altered for holidays, with prior notice to the Review Panel.

k) The remaining days of the review cycle are for the submitter to resolve any issues with the result under review.

   Notes:

   - Minor changes or clarifications will be allowed in the current review cycle.
   - A rerun of a result to make a minor correction is allowed, but neither component of the rerun score (the cumulative work or the number of tiles) can be higher than the corresponding component of the original score.
   - Any changes to the testing environment not required to address a compliance issue will require a re-review in a subsequent review cycle.

l) At the end of the review cycle, a result will be accepted by VMware for publication on the VMmark web page if either no issues were flagged, or if all flagged issues were satisfactorily addressed prior to the end of the review cycle.

m) VMware will notify the submitter with the final status of the review for each separate submittal.

n) It is expected that under normal circumstances an accepted result will be posted to the VMmark web page within 24 hours of the last day of the review cycle. Submitters can also request a slightly delayed publication date if desired.

If it is determined that the VMmark test was not run in full compliance with these rules, VMware will not accept that result for publication and the benchmarker must not use that result in any

public disclosure. The benchmarker has the option of rerunning the test after making changes needed to bring it into compliance with these rules and submitting the new result and full disclosure report for review. Any such rerun must be uploaded before 9:00 am Pacific time two weeks after the close of the review cycle for which the submission was received.

The benchmarker may withdraw a submission at any time while it is under review. However, any open issues must be addressed prior to publishing the result independently.

2.2 Independent Publication of VMmark Results

Benchmarkers can independently publish results from VMware approved fully compliant tests for which a full disclosure report (as described in Section 2.3) is publicly available. After independent publication the benchmarker can, if they wish, formally submit the full disclosure report for publication on the VMmark website. Benchmarkers cannot, however, independently publish a result while it is in a review panel review cycle. Any compliance issues found by VMware or other readers of an independently published disclosure will be handled in the same manner as results that have undergone the formal review process.

2.3 Submission and Disclosure Requirements

VMmark result submissions must include a full disclosure report (FDR) of the SUT (as defined in Section 1.1), the datacenter management server, the clients, and all configuration and tuning information. In the case of VMmark results that include power, the FDR must also include all relevant power information. The FDR must be sufficiently detailed that a third party could use this information to duplicate the test and obtain results within a deviation of +/- 3%. The level of detail in the FDR must be, at a minimum, comparable to other industry standard benchmark results. The full set of VMmark test result measurements and metrics must also be included. All result submissions must use the correct FDR HTML template. This will be the template included with the latest VMmark release available at the time of submission, unless multiple VMmark releases are available, in which case results submissions must use the latest FDR HTML template for the VMmark version being submitted. For example, if both VMmark 3.0 and 3.1 releases are available, a VMmark 3.0 submission must use the latest VMmark 3.0 FDR HTML template.

A full disclosure report must be available from either VMware or the benchmarker whenever a VMmark result is published. The disclosure report can either be included as a hypertext link (URL) or included in its entirety in the publication.

The following information must accompany any VMmark result that is submitted for review or, in the case of independently published results, be made available on request:

a) The contents of the prime client's results directory (VMmark3/results/Results_<datestamp>), archived in .zip or .tar.gz format. This directory will contain the workload result files (.wrf) for all workloads and all tiles run.

b) The output from any scripts supplied with the harness reporting tools to collect configuration details on the benchmark environment, archived in .zip or .tar.gz format. Alternately, if the supplied scripts do not support a given virtualization environment, configuration details can be collected using individual commands or supplemental software provided by VMware.

A single benchmark run may be used for one or more VMmark submission types (i.e., VMmark Performance, VMmark Performance with Server Power, or VMmark Performance with Server and Storage Power). A separate FDR is required for each type of VMmark result being submitted. For example, the submitter may submit all of the data from a benchmark run, with a separate FDR for one, two, or three types of results.

2.4 VMmark Scores

There are two types of VMmark scores: Performance scores (from Performance Only results) and PPKW scores (from Server Power-Performance and Server and Storage Power-Performance results), each of which consist of two components.

- Performance scores are of the form X @ Y tiles. The "X" component of the score represents the cumulative work performed by all the workloads in the test. The "Y tiles" component of the score indicates how many VMmark tiles were used to produce the result.
- PPKW scores are of the form X @ Y tiles. The "X" component of the score represents the cumulative work performed by all the workloads in the test divided by the average power (in kilowatts) consumed during the test. For the Server Power-Performance results, this is the average power consumed by the servers. For the Server and Storage Power-Performance results, this is the average power consumed by both the servers and the storage. The "Y tiles" component of the score indicates how many VMmark tiles were used to produce the result.

2.5 Comparability of VMmark Results

a) VMmark 3.x defines three result types:

   - VMmark Performance
   - VMmark Server Power-Performance
   - VMmark Server and Storage Power-Performance

   Results of any one type are strictly not comparable to results of another type.

b) Only official VMmark metrics and submetrics may be used in comparisons.

c) Comparisons of VMmark metrics and submetrics to any other benchmark metrics are not allowed. VMmark utilizes other benchmarking software as load generators and produces results that are not comparable to the original benchmarks' metrics.

d) Competitive comparisons are those comparisons that involve VMmark results where one or more performance-critical components or products from competing vendors have been used.

e) In a competitive comparison, the VMmark result upon which a claim or implication of superiority is based is defined as the "advocated result." The claim or implication needn't be that the advocated result has a higher VMmark score; it can include other factors such as cost, power consumption, and so on. For example, if a vendor claims that their system has higher performance per dollar because, even though its VMmark score is half that of another system, it is only one third the cost, then that vendor's VMmark result is an advocated result.

f) VMmark results are not comparable if they were produced using VMmark releases with different major release numbers. Results with different major release numbers are grouped separately on the VMmark website. Notifications will be posted when a new release is available and will indicate the last date results based on the prior release will be accepted. VMmark 1.x and 2.x results are not comparable to VMmark 3.x results.

g) Competitive comparisons between two VMmark results produced with the same major release number can be made, but only with the following limitation: All security vulnerabilities that were required by the VMmark Run Rules and the *VMware VMmark User's Guide* to be mitigated in the non-advocated result must also have been mitigated in the advocated result. There is no requirement, however, that the mitigations be accomplished using the same method, nor is the advocated result prohibited from including additional mitigations.

h) The power measurement capability in VMmark 3.x utilizes the SPEC PTDaemon (Power Temperature Daemon). VMmark results are not SPEC metrics and cannot in any manner be compared to SPEC metrics.

2.6 Handling Compliance Issues in Published Results

VMware and the benchmarker will promptly address any issues related to compliance with these rules that are raised after the result has been published. If such issues are raised, VMware may determine that the result should be withdrawn. In this case both parties must withdraw any web pages or downloadable reports that reference this result and must stop using the result in printed or other media. The materials can be revised to remove the reference and republished.

VMware will maintain a “Withdrawn Results” web page that provides a summary of the tested platform and the reason the published report was designated either non-compliant (NC), not available (NA), or code defect (CD) and was withdrawn from the “Current Results” web page. The benchmarker may request that the entry be updated to include the description of any remedial action they've taken. For example, if the reason the result was designated non-compliant was: “Failed to meet 90-day Public Availability requirement for products used,” it later could be updated to say: “All products used were publicly available as of 06/07/08 and a new result has been submitted.”

2.7 Handling of Previously Published Results

VMware retains the right to publish and reference all results that have been reviewed by the Review Panel and accepted by VMware for publication. Once a result is published on the VMmark results page, it will remain there unless it is later found to be non-compliant, not available, or code defect (in which case it will be moved to the "Withdrawn Results" page, as described in Section 2.6, above).

2.8 Fair Use: Handling of Published VMmark References

VMware requires that all published references to VMmark results (except those published under the academic and research rules) adhere to the following fair-use rules:

a) All VMmark results referenced must be from fully compliant tests for which a full disclosure is publicly available, either through the VMmark website or through a site specified by the benchmarker. The use of estimates is prohibited.

b) A reference or link to the location of the full disclosure report for all referenced results must be included.

c) Comparisons of results are allowed only according to the definition given in Section 2.4 of this document. If comparisons are made between “current” and “historical” results with the same major release number, the “historical” result must be clearly identified as “historical.”

d) Any competitive comparisons of VMmark results must include a statement providing the basis for making the comparison and the date when the claim is first made. A published comparison could, for example, state that the claim is based on having the best VMmark score out of all four-socket results published as of the publication date (with a specific MM/DD/YYYY date given).

e) All publications that reference VMmark results must include the following attribution: VMmark® is a product of VMware, Inc.

2.9 Academic and Research Publication of VMmark Results

In addition to publication on the VMmark results page (as described in Section 2.1) and independent publication (as described in Section 2.2), VMware permits academic and research use of VMmark, described in this section.

2.9.1 Limitations for Academic and Research Publications

Publication of VMmark results under the academic and research rules is subject to the following conditions:

Note: Publication refers to disseminating VMmark results in a public setting. Publication under academic use policy could include, but is not limited to, scholarly articles, whitepapers, blog articles, books, and online references.

a) All such publications must be clearly identified as "academic" or "research."

b) Disclosure of all variations from fully compliant test runs must be included in the publication.

c) Such published results must not utilize the VMmark metric names ("VMmark Applications Score," "VMmark Infrastructure Score," "VMmark Score," and "VMmark PPKW").

d) Instead of VMmark metric names, such published results should use normalized metrics. For example, a result might state that a virtualization environment performed 25% better with an experimental tuning than without it.

e) Competitive comparisons (as defined in Section 2.7) using VMmark results in academic or research papers are not allowed.

f) Even though academic or research results don't require pre-publication review by VMware under the VMmark rules, such results normally require pre-publication review by VMware under the general VMware benchmarking approval process.
Note: The general VMware benchmarking approval process is waived if the university is part of the VMware Academic Program (VMAP). For more information about the VMAP see https://labs.vmware.com/academic/licensing-overview.

2.9.2 Exceptions for Academic and Research Publications

Publication of academic or research VMmark results differs from publication on the VMmark results page or independent publication in a number of ways.

Academic and research publications:

a) Do not require a Full Disclosure Report.

b) Do not require formal submission (that is, review by the VMmark Review Panel) prior to publication.

c) Do not require review by VMware prior to publication unless the publication falls under the general VMware benchmarking approval process.
Note: The general VMware benchmarking approval process is waived if the university is part of the VMware Academic Program (VMAP). For more information about the VMAP see https://labs.vmware.com/academic/licensing-overview.

d) Do not require a compliant VMmark setup (that is, a three-hour run, all workloads active, two or more hosts, and following the VMmark Run Rules for a compliant result).

e) Can make use of hardware and software components that are not Publicly Available at the time of publication or within 90 days after publication. In this case, the associated report must clearly identify the components that are not Publicly Available and their planned release dates or indicate that the component was experimental and not planned for product release.

f) Can use Beta versions of VMmark, provided such results:
- Are clearly marked "Beta VMmark Results,"
- Include the Version ID from the Version.txt file included in the release,
- Were performed using released and currently supported versions of VMware vSphere and VMware ESXi.

Note: Beta versions of VMmark may be used only until a final release version becomes available.

3 VMmark Run Rules

The VMmark Run Rules describe the requirements for configuring the benchmark environment to produce valid and compliant test results. This section references the *VMware VMmark User's Guide* (the *"User's Guide")*, which provides specific requirements for configuration of the individual workload virtual machines (VMs), the clients, the test environment, and (optionally) power meter set up.

In addition to the requirements included in these Run Rules, the required virtual hardware definitions, the required software versions for the workload VMs, and all other requirements included in latest version of the *User's Guide* must be followed. Any question regarding compliance with either document should be clarified with VMware to ensure the creation of valid test results.

3.1 Overview *of the VMmark Benchmark*

VMmark 3.x is a multi-host virtualized datacenter benchmark. A tile is defined as 19 VMs running the workloads specified in the *User's Guide.* Tile VMs are run across multiple servers while the benchmark simultaneously performs periodic infrastructure operations.

Additionally, VMmark allows the measurement of power usage while running its full application and infrastructure workloads across a multi-host cluster by providing for three test types:

- Performance only
- Performance with Server Power
- Performance with Server and Storage Power

3.1.1 VMmark 3.x Benchmark Workloads

A VMmark tile consists of a collection of workload virtual machines and an associated client virtual machine.

The instructions to create a VMmark tile are included in the *User's Guide*. Those instructions must be followed in order to create a compliant VMmark tile. Changes made to any of the virtual machines that make up the tile, other than those specifically allowed in these Run and Reporting Rules or in the *User's Guide*, potentially make non-compliant any benchmark run containing that tile, and any resulting submission.

3.2 SUT Configuration Requirements

This section covers configuration requirements for the SUT.

3.2.1 Performance Neutrality

On rare occasions, a VMmark result obtained on a SUT containing a configuration, tuning, version, or other factor not permitted under these Run and Reporting Rules might nevertheless be accepted as compliant if VMware deems the factor in question to be performance neutral. Typically a performance neutral factor is one that will have no effect on the default settings as tested (e.g., the factor is ignored or is overridden by another setting) or will have such little effect on performance that the result could be duplicated without this factor (e.g., enabling a virtual device not used in the test).

Note that in addition to performance neutrality, other considerations enter into the decision to accept as compliant a result with a non-permitted factor; thus, even a clearly performance neutral factor could keep a VMmark result from being accepted as compliant. If you have questions or concerns about performance neutrality, you are strongly encouraged to contact VMware for clarification.

3.2.2 vSAN Storage Results

A VMmark result will be considered a vSAN Storage result, and will be presented as such on the VMmark results page, if all application workload virtual machines in the result are placed on VMware vSAN storage. Use of non-vSAN storage hardware for infrastructure target storage (that is, the storage used for the Storage vMotion, Cross-host Storage vMotion, and deploy targets) and for the deploy template does not disqualify a result as a vSAN Storage result.

3.2.3 Uniform Versus Non-Uniform Hosts

A VMmark result can be submitted in either the "uniform hosts" or "non-uniform hosts" categories. The following sections describe these two categories and the requirements for each one.

3.2.3.1 Uniform Hosts

A VMmark result will be considered a uniform host result if it is obtained on a SUT within which all the compute hosts are identical (that is, are the same brand and top-level model) and the initial disclosure declares it a uniform host result. Uniform host results obtained on SUTs containing exactly two compute hosts will be further listed as being "Matched Pair" results.

In uniform host results, the compute hosts within the SUT must have identical configurations (including hardware, firmware, and software), and the compute hosts must all run the same hypervisor version. On rare occasions, variations might be allowed as performance neutral, as described in Section 3.2.1.

3.2.3.2 Non-Uniform Hosts

A VMmark result will be considered a non-uniform host result if:

a) It is obtained on a SUT containing compute hosts that are not identical (that is, are not the same brand and top-level model) or are identical but configured or tuned differently and

b) The initial disclosure:

- declares it a non-uniform host result,
- summarizes the non-uniformity in the "Vendor and Hardware Platform" portion of the title,
- describes in the Configuration section what makes the compute hosts non-uniform, and
- describes in the Notes section what the test is designed to show and why non-uniform compute hosts are appropriate for that purpose.

Acceptance of a VMmark submission as a non-uniform host result is at VMware's discretion. Before performing a VMmark test intended for submission as a non-uniform host result, we encourage you to contact VMware with a description of your SUT and the reason you plan to use non-uniform hosts. While prior approval is not required, failing to consult with VMware before submitting a non-uniform host result could delay the submission review process or even result in the submission being deemed non-compliant for reasons that might have been easily addressed in advance.

In non-uniform host results, the compute hosts within the SUT all run the same hypervisor version. On rare occasions, variations might be allowed as performance neutral, as described in Section 3.2.1.

3.2.4 Versions, Tunings, and Configurations

The benchmarker must configure their VMs according to the instructions in the *User's Guide* and follow the limitations on the software versions as described. The benchmarker may apply updates to these releases as described in the *User's Guide* but may not use other versions of Linux or substitute one OS for another.

The server, client, network, and power monitor hardware configuration requirements, as well as the associated software requirements, are also detailed in the "VMmark Benchmark Requirements" section of the *User's Guide*. To produce fully compliant VMmark results, all components in the benchmark environment must have been configured according to these requirements.

The security mitigations described in the *User's Guide* must be enabled.

The benchmarking environment must be configured such that there are no DRS-induced vMotion failures.

Certain vSphere tunings are permitted, as detailed in Section 3.2.5.2. The permitted tunings are vSphere features that a datacenter manager might use to optimize their datacenter, while at the same time not altering the fundamental structure and fairness of the underlying benchmark methodology.

3.2.5 Types of Changes Permitted to the SUT

This section describes the types of changes to the SUT that are generally permitted by these VMmark Run and Reporting Rules and the *User's Guide*.

Note: In addition to the changes described in this section, a change or tuning might be permitted if it is deemed performance neutral by VMware, as described in Section 3.2.1.

3.2.5.1 SUT Hardware Configuration

a) The hardware configurations for the SUT may be modified or tuned as long as they meet the support requirements for the hardware vendor(s) and the vendor of the virtualization product being used. Changes may include adding memory, NICs, or storage adapters, attaching storage arrays, etc.

b) Changes to the server's BIOS or firmware parameters may be made if consistent with the vendors' published recommendations.

c) All storage and file systems must use stable non-volatile media.

3.2.5.2 SUT Virtualization Layer Configuration

a) The virtualization layer on the SUT can be modified or tuned as long as it remains a supported configuration and the tuning is consistent with the vendor's published recommendations.

b) VMmark permits certain vSphere tunings. Other general tuning is allowed, but any other tunings that alter the resource allocation to favor some VMs over others are not permitted. The defined list of permitted tunes for resource control are:

    a. Setting CPU shares

    b. Setting memory shares

    c. Setting memory reservations

       Note: Memory reservations may be used only as long as every server in the SUT is configured to have a minimum amount of free memory equaling at least 16GB + 2% of the total memory installed in the server, as reported by the VMmark harness, over the entire measurement interval of the run.

    d. vmkernel vSwitch and VDS traffic shaping

    e. Using Storage I/O Control to set congestion thresholds

    f. Using Storage I/O Control to set disk shares

c) Any tunings must be applied to all VMs of the same workload type across all tiles used in the test (unless the tuning is deemed by VMware to be performance neutral for this SUT, as described in Section 3.2.1). For example, if a virtualization optimization is enabled on the Weathervane Auction DB VM in one tile, then it must be enabled on the Weathervane Auction DB VMs in every tile. Tunings can be achieved either as a per-VM configuration, or as a resource pool, provided that if any VM of a specific workload type is part of a resource pool, all other VMs of that type are part of the same resource pool.

3.2.5.3 Workload Virtual Machine Configuration

a) All workload VMs must use the VMmark template supplied by VMware. The platform definition – which includes the number of vCPUs, the amount of memory, the virtual hardware, the storage and network requirements, and the software requirement for each workload VM – must not be changed, except as specified in the *User's Guide*.

b) The workload VMs' virtual hardware version can be changed to a later version when specified in the *User's Guide*.

c) The workload VMs' VMware Tools version can be changed to a later version when specified in the *User's Guide*.

d) Each workload VM must be configured with only one virtual NIC with the default Maximum Transmission Unit (MTU) of 1500 bytes.

e) The version of the guest OS or the guest OS kernel that is included with the VMmark template or an allowed update release must be used.

f) The default system services and daemons must remain enabled unless specific changes are specified in the *User's Guide*.

g) If any workload VM's software configuration or virtual hardware version is changed based on the rules above (such as upgrading the VM's OS to an allowed update release), then that change must be applied to all VMs of that type in all tiles (unless the tuning is deemed by VMware to be performance neutral for this SUT, as described in Section 3.2.1), and information about the change must be listed in the Notes section of the disclosure.

h) Software updates supplied by VMware to address problems or supply new features may be used in place of VMmark components in the current release. Such changes must be documented in the notes section of the disclosure.

3.2.6 Types of Changes NOT Permitted to SUT

The following section describes the types of changes to the SUT that are generally NOT permitted by these VMmark Run and Reporting Rules and the *User's Guide*. Refer to the *User's Guide* for additional detail, including specific parameter changes that are prohibited.

3.2.6.1 SUT Hardware Configuration

Features that are not consistent with the vendors' published recommendations may not be used in the SUT.

3.2.6.2 SUT Virtualization Layer Configuration

a) Features that are not consistent with the vendors' published recommendations may not be used in the configuration of the virtualization layer.

b) Features that are not consistent with the vendors' published recommendations may not be used in the per-VM virtualization configuration.

c) Beyond the defined list of permitted tunings for resource control given in section 3.2.5.2, configuration changes that alter the balance of assigned resources (CPU, memory, disk, network) for specific VMs are not permitted. This includes optimizations that increase or decrease the amount of a given resource to be allotted to a specific VM in the tile when the system is under load.

3.2.6.3 Workload Virtual Machine Configuration

a) No changes may be made to the VMmark template supplied by VMware except as explicitly permitted in these Run and Reporting Rules and in the *User's Guide*.

b) No changes may be made to the tuning parameters for the guest OS unless specified in the *User's Guide*.

c) No changes may be made to any workload-specific application or benchmarking code that runs on the workload VMs unless the *User's Guide* includes specific directions to do so.

d) No changes may be made that alter the resources allotted to specific workload VMs to the advantage or disadvantage of any other VM in the tile.

e) No changes may be made to any code supplied as part of the VMmark distribution unless the *User's Guide* specifies that the file or designated parameters can be changed.

3.2.7 Types of Changes to SUT Not Covered

Any changes, which may or may not have been explicitly allowed or disallowed by these rules or the *User's Guide,* should be reviewed with VMware prior to their use.

3.2.8 SUT Product Availability and Support Requirements

All hardware and software components that compose the SUT must meet the availability and support requirements detailed in this section.

a) SUT availability requirements are separated into three categories: server, storage, and infrastructure target storage:

   i. **Server:** Unless stated otherwise in this section, the hardware and software components that compose the SUT must either be Publicly Available or be pre-release versions within 90 days of official Public Availability.

      Note: If a component of the SUT is not specifically mentioned in other parts of this section, it should be considered to be part of the server category.

   ii. **Storage:** This section addresses storage components used in the SUT. This category is defined as disks and storage arrays and their related components (including NAS and Fibre Channel networking components), but not hyperconverged infrastructure (HCI) storage, such as vSAN (which is considered part of the server for purposes of availability and support). The storage components defined here must either have been Publicly Available at some time within the last five years or be pre-release versions within 90 days of official Public Availability.

      Note: This means that storage components are not required to have still been available for purchase at the time of the benchmark run, nor are they required to have still been supported at that time. However, the disclosure must state the original Public Availability date, the date the product became unavailable, and the date support became unavailable.

   iii. **Infrastructure target storage:** If infrastructure target storage (that is, the storage used for the Storage vMotion, Cross-host Storage vMotion, and deploy targets) is not located on the same storage hardware as the application workload VMs, its components must either have been Publicly Available at some time within the last five years or be pre-release versions within 90 days of official Public Availability. These components must be (or have been) commercially-supported, but they don't need to have ever been supported for vSphere.

      Note: This means that infrastructure target storage components are not required to have still been available for purchase at the time of the benchmark run, they are not

required to have still been supported at that time, and they are not required to have ever been officially supported for use with vSphere. However, the disclosure must state the original Public Availability date, the date the product became unavailable, and the date support became unavailable.

b) After first becoming Publicly Available, all products used must remain on the market for 60 of the first 90 days and must remain fully supported for one year. If a product is removed from the market for more than 30 days during the first 90 days, or if it stops being fully supported during the first year, it is considered "Not Available" (NA). Replacement of components within the SUT by equivalent components or new versions with equal or better performance is considered acceptable and not a reason for a NA designation.

c) During the first 90 days after the SUT becomes Publicly Available or after the VMmark result is published (whichever is later), the products that comprise the SUT must continue to offer performance equivalent to or better than their performance in the VMmark result. To have equivalent performance, any regression must reduce the VMmark result by no more than 3%.

d) To preserve the integrity of the benchmark and the software components it uses, open source software ("OSS") that is not commercially distributed and supported must not be used in the SUT unless the OSS meets certain requirements. The requirements, which are evaluated based upon the opening date of the review cycle into which the benchmark run is submitted, include:

   i. The OSS must be licensed under one of the OSI-approved licenses listed at http://www.opensource.org/licenses/ or another license with prior approval from VMware.

   ii. The OSS must be a release that is:
      A. a *stable* release, requiring that the OSS is not alpha, beta, or other type of pre-release;
      B. an *established* release, requiring that:
         - the first stable release of the OSS is more than one year old; and,
         - at least two additional stable releases of the OSS have been released (following the first stable release);
      C. a *recent* release, requiring that:
         - the release was the latest stable release at some point within the previous six months; and,
         - the release is less than 18 months old;
      D. an *active* release, requiring that at least one identified developer is actively contributing to and maintaining the OSS.

      Note: In rare instances, VMware might deem the OSS a *mature* release and grant an exception to the *recent* release and *active* release requirements. When such an exception is granted, it will typically be because:
         - the OSS requires little or no maintenance; and,
         - there might no longer be active development of the OSS.

      Benchmarkers are encouraged to request such an exception in advance.

   iii. The OSS must have an online support forum ("OSF") that is:

A. an *active* OSF, requiring that:
    - the OSF includes at least one posting authored within the last three months; and,
    - such posting was not authored by the benchmarkers, their representatives, or members of the VMmark review panel;
B. a *helpful* OSF, requiring that:
    - the OSF provides useful responses to users' questions, such that if a previously unreported problem is reported with sufficient detail, it is responded to by a project developer or community member with sufficient information that the user ends up with a solution, a workaround, has been notified that the issue will be addressed in a future release, or has been told that it is outside the scope of the project; and,
    - the OSF archive includes at least one instance of this level of conversation;
C. a *functional* OSF, requiring that:
    - the OSF is easily accessible and usable;
    - the OSF does not place restrictions on access; and,
    - the OSF provides access to all previous postings.

3.3 Client Configuration Requirements

This section covers configuration requirements for the client.

3.3.1 Types of Changes Permitted to Clients

This section describes the types of changes to clients that are generally permitted by these VMmark Run and Reporting Rules and the requirements sections of the *User's Guide*.

Note: In addition to the changes described in this section, a change or tuning might be permitted if it is deemed performance neutral by VMware, as described in Section 3.2.1.

3.3.1.1 Client Virtualization Layer Configuration

d) The virtualization layer can be modified or tuned as long as the tuning is consistent with the vendor's published recommendations.

e) VMmark permits certain vSphere tunings. Other general tuning is allowed, but any other tunings that alter the resource allocation to favor some VMs over others are not permitted. The defined list of permitted tunes for resource control are:

   a. Setting CPU shares
   b. Setting memory shares
   c. Setting memory reservations
   d. vmkernel vSwitch and VDS traffic shaping
   e. Using Storage I/O Control to set congestion thresholds
   f. Using Storage I/O Control to set disk shares

f) Any tunings must be applied to all the tile client VMs across all tiles used in the test (unless the tuning is deemed by VMware to be performance neutral, as described in Section 3.2.1). For example, if a virtualization optimization is enabled for the tile 0 client VM, then it must be enabled for all the other tile client VMs. Tunings can be achieved either as a per-VM configuration, or as a resource pool, provided that all the tile clients are part of the same resource pool.

3.3.1.2 Client Virtual Machine Configuration

a) All client VMs must use the VMmark template supplied by VMware. The platform definition – which includes the number of vCPUs, the amount of memory, the storage and network requirements, and the software requirement for each workload VM – must not be reduced. Virtual hardware resources can be increased, however, unless otherwise specified.

b) The virtual hardware version used for the client VMs can be changed to a later version when specified in the *User's Guide*.

c) Each client VM (with the exception of the prime client) must be configured with only one virtual NIC with the default Maximum Transmission Unit (MTU) of 1500 bytes.

d) The version of the guest OS and guest OS kernel that is included with the VMmark template or an allowed update release must be used.

e) The default system services and daemons must remain enabled unless specific changes are specified in the *User's Guide*.

f) If any tile client VM's software configuration is changed based on the rules above (such as upgrading the VM's OS to an allowed update release), then that change must be applied to all tile client VMs (unless the tuning is deemed by VMware to be performance neutral for this SUT, as described in Section 3.2.1), and information about the change must be listed in the Notes section of the disclosure.

3.3.1.3 VMmark Benchmark Harness Configuration

a) The VMmark Benchmark Harness software installed includes the VMmark3.properties configuration file, which may be modified as described in the *User's Guide*.

b) Software updates supplied by VMware to address problems or supply new features may be used in place of VMmark components in the current release. Such changes must be documented in the notes section of the disclosure.

3.3.2 Types of Changes NOT Permitted to Clients

The following section describes the types of changes to the clients that are generally NOT permitted by these VMmark Run and Reporting Rules and the requirements sections of the *User's Guide.*

3.3.2.1 Client Hardware Configuration

Features that are not consistent with the hardware vendors' published recommendations may not be used in the client systems.

3.3.2.2 Client Virtualization Layer Configuration

d) Features that are not consistent with the virtualization layer vendors' published recommendations may not be used in the configuration of the virtualization layer.

e) Features that are not consistent with the virtualization layer vendors' published recommendations may not be used in the per-VM virtualization configuration.

f) Beyond the defined list of permitted tunings for resource control given in section 3.2.5.2, configuration changes that alter the balance of assigned resources (CPU, memory, disk, network) for specific tile client VMs are not permitted. This includes optimizations that increase or decrease the amount of a given resource to be allotted to a specific tile when the system is under load.

3.3.2.3 Client Virtual Machine Configuration

f) No changes may be made to the VMmark template supplied by VMware except as explicitly permitted in these Run and Reporting Rules and in the *User's Guide*.

g) No changes may be made to the tuning parameters for the guest OS unless specified in the *User's Guide*.

h) No changes may be made to any code that runs on the client VMs unless the *User's Guide* includes specific directions to do so.

i) No changes may be made that alter the resources allotted to specific VMs to the advantage or

disadvantage of any other VM in the tile.

3.3.2.4 Client Code and Other Configuration

a) No changes may be made to any workload-specific client-side code unless the *User's Guide* includes specific directions to do so.

b) No changes may be made to any code supplied as part of the VMmark distribution unless the *User's Guide* specifies that the file or designated parameters can be changed.

3.3.2.5 VMmark Benchmark Harness Configuration

a) The VMmark Benchmark Harness configuration files may not be modified except as described in the *User's Guide*.

b) Changes to the harness software are not allowed with the exception of updates supplied by VMware.

3.3.3 Types of Changes to Clients Not Covered

Any changes, which may or may not have been explicitly allowed or disallowed by these rules or the *User's Guide,* should be reviewed with VMware prior to their use.

3.3.4 Client Product Availability and Support Requirements

All hardware and software components that compose the clients must meet the availability and support requirements detailed in this section.

a) The hardware and software components either are Publicly Available products, were at some time in the past Publicly Available products, or are pre-release versions within 90 days of official Public Availability.

   Note: This means that client components are not required to have still been available for purchase at the time of the benchmark run and they are not required to have still been supported at that time.

b) All products used must remain on the market for 60 of the first 90 days after they first become Publicly Available. If a product is removed from the market for more than 30 days during the first 90 days, it is considered "Not Available" (NA). Replacement by a new similar component or new version with equivalent or better performance is considered acceptable and not a reason for a NA designation.

4 Handling Run and Reporting Rule Compliance

To ensure that benchmarkers adhere to the VMmark Run and Reporting Rules, this section summarizes typical compliance issues and how they will be handled to ensure the integrity of the benchmark and VMmark metrics.

4.1 Compliance Issues Found Prior to Publication

While a submission is in review, issues related to the full disclosure report (FDR) will be brought to the attention of the submitter so that the report can be updated. If the issues raised point to a possible violation of the run rules, VMware might ask the submitter to clarify the situation, to perform additional runs using the same SUT and configuration, or to rerun the benchmark in compliance with the run rules. The submitter may withdraw the submission at any time during the review period but may not independently publish the result if there are open compliance issues.

4.2 Compliance Issues after Publication on VMmark Website

Any compliance issues found after a result is published will be brought to the attention of the VMmark Review Panel. If a compliance issue in the FDR is found after the submission has been published on the VMmark website, the submitter may correct the FDR and it will be published in place on the VMmark Results page. If the compliance issue raised points to a run rule violation and further review by the Review Panel substantiates this, then VMware will designate the published result as Non-Compliant and update the VMmark website as described in Section 2.6. The Submitter is encouraged to make a new compliant run and submit the new result for review and publication.

4.3 Compliance Issues After Independent Publication

If a compliance issue with the FDR is found after the submission has been published outside the VMmark website, then VMware will work with the benchmarker to correct the FDR for re-publication. If the compliance issue raised points to a run rule violation, and further review by the Review Panel substantiates this, then VMware will request that references to the result in question be taken down (see Section 2.6). The benchmarker is encouraged to make a new compliant run and republish as appropriate. VMware reserves the right to take action if VMmark results that have been designated Non-Compliant remain posted on the submitter's website or other media which they control.

4.4 Compliance Issues When a Pre-Release Product Is Not Released

If pre-release products have been used and those products have not become Publicly Available by the end of the 90-day window (see Sections 3.2.8 and 3.3.4) then the result will be designated Non-Compliant (NC). If the result has been published on the VMmark website, then VMware will designate the published result as Non-Compliant and update the VMmark website as described in Sections 2.6 and 4.2 above. Any publication of the result outside the VMmark website will be handled as described in Sections 2.6 and 4.3 above.

4.5 Compliance Issues When a Product Is Not Available

If it is found that within 90 days of the publication of a VMmark result or the Public Availability

date of the SUT, whichever is later, performance critical products used in the SUT had been taken off the market for more than 30 days during that period, then that result will be designated Not Available (NA). If the result has been published on the VMmark website, then VMware will designate the published result as Not Available and update the VMmark website as described in Section 2.6. Any publication of the result outside the VMmark website will be handled as described in Section 2.6. The benchmarker is encouraged to make a new compliant run and submit the new result for review and publication. Note: Replacement by a new similar component or new version with equivalent or better performance is considered acceptable and not a reason for an NA designation.

4.6 Compliance Issues When a Result Is Not Reproducible

If within 90 days of the publication of a VMmark result, or the Public Availability date of the SUT, whichever is later, it is found that there has been a performance regression of more than 3%, then that result will be designated non-compliant. If the result has been published on the VMmark website, then VMware will designate the published result as Non-Compliant and update the VMmark website as described in Sections 2.6 and 4.2 above. Any publication of the result outside the VMmark website will be handled as described in Sections 2.6 and 4.3 above.

4.7 Compliance Issues Due to a Code Defect

If a published result is found to be non-compliant due to a code defect in the VMmark harness or any of its supporting scripts, then VMware will designate the published result as Code Defect (CD) and update the VMmark website as described in Sections 2.6 and 4.2 above. Any publication of the result outside the VMmark website will be handled as described in Sections 2.6 and 4.3 above.

5 Caveats

VMware reserves the right to modify the VMmark benchmark codes, workloads, and documentation, including these Run and Reporting Rules, as needed to ensure the value and integrity of VMmark as a performance benchmark for virtualization. The latest published edition of these Run and Reporting Rules will apply to all VMmark 3.x releases and VMmark 3.x results unless specifically stated to the contrary.

If VMmark users have questions or concerns regarding these rules and their interpretation, or about any of the VMmark materials, they are strongly encouraged to contact VMware for clarification. Additional feedback on VMmark is welcome as well. VMware provides a VMmark forum that can be used to post questions and provide feedback (see: communities.vmware.com/community/vmtn/server/vmmark).