

Guide to Vulkan Synchronization Validation

John Zulauf, LunarG

Version 2.1, released with SDK 1.3.275.0

February, 2024

Table of Contents

Table of Contents	2
Introduction	2
Quick Start	2
Using Synchronization Validation	2
Synchronization Validation Messages	3
Frequently Found Issues	3
Debugging Tips	4
What is synchronization and why is it important?	4
Synchronization Operations	6
Pointers to synchronization blogs/articles	7
Using Synchronization Validation	7
Enabling synchronization validation.	7
Using Vulkan Configurator	7
Using Layer Settings	7
Using Environment Variables	8
Typical Usage	8
Understanding Synchronization Validation Messages	9
Message Contents	9
Recorded/Executed/Submitted Comparative Example	11
Understanding Synchronization Validation	13
Stage/Access Usage Pairs	13
Hazard Detection	14
Most Recent Access	14
Read Operations	14
Write Operations	14
Debugging Hazards	17
Debugging Using Access info information	19
Hazards from Missing or Incomplete Barriers	19
Zero (empty) Read and Write Barriers	19
Non-Zero Barriers	19
Hazards vs. Prior Image Layout Transitions	20
Hazards at Image Layout Transitions	20
Hazards between buffer and/or image resource uses	20
Method of Bisection Using Additional Barriers	21
Identifying Affected Resources and Operations	21
Getting Consistent Resource Identification	21

Tracking Operations For a Given Resource	22
Using Code Inspection	22
Look near the stack trace location	22
Identifying Incomplete Existing Barriers	22
Examining Resource Use Transitions	23
Optimizing Synchronization with Synchronization Validation	23
For Further Information and Current Status	24
Revision history	24

Introduction

Synchronization Validation is implemented in the `VK_LAYER_KHRONOS_validation` layer as optional additional validation. When enabled, it is intended to identify resource access conflicts (often referred to as “hazards”) due to missing or incorrect synchronization operations between actions (Draw, Copy, Dispatch, Blit) reading or writing the same regions of memory. Synchronization Validation also reports issues between or outside of recorded command buffers (Queue Submit, etc.).

Quick Start

The following quick start should enable initial testing for those familiar with Vulkan synchronization and debugging validation issues. Prior to enabling Synchronization Validation, assure that the default set of Validation checks run cleanly.

Using Synchronization Validation

The simplest way to run synchronization validation is to enable Synchronization Validation using [Vulkan Configurator \(vkconfig\)](#). This will report synchronization issues with “HAZARD” error messages. (See “Understanding Synchronization Validation Messages” below.)

Debugging reported “HAZARD” errors is more complex, is best done in an interactive debugger, and may involve some code modification. Typical steps include:

- Resolve all validation error messages from the default Vulkan Validation configuration. Invalid Vulkan usage will affect the correctness of Synchronization results.
- Enable Synchronization Validation using [Vulkan Configurator \(vkconfig\)](#).
- Create a debug callback with `vkCreateDebugUtilsMessengerEXT` with `VK_DEBUG_REPORT_ERROR_BIT_EXT` set.
- Set a breakpoint in the debug callback and run your application in the debugger.
- The hazards will be reported when a `vkCmd...` command with a hazard is recorded, or when a command buffer containing a `vkCmd...` command with a hazard is submitted.

Synchronization Validation Messages

Synchronization messages report conflicts between memory accesses in the current Vulkan API command and memory accesses *prior* commands. All synchronization error messages begin with `SYNC-<hazard name>`. The message body is constructed:

```
<cmd name>: Hazard <hazard name> <command specific details> Access info (<...>)
```

Access info contains information about current and prior usage (formatted `SYNC_<stage>_<access>`) and any intervening synchronization. Memory or subresource range of the usage is given in the command-specific details among other information.

Frequently Found Issues

- Assuming pipeline stages are logically extended with respect to memory access barriers. Specifying the vertex shader stage in a barrier will **not** apply to all subsequent shader stages read/write access.
- Invalid stage/access pairs (specifying a pipeline stage for which a given access is not valid) that yield no barrier.
- Relying on implicit subpass dependencies with `VK_SUBPASS_EXTERNAL` when memory barriers are needed.
- Missing memory dependencies with Image Layout Transitions from pipeline barrier or renderpass Begin/Next/End operations.
- Missing stage/access scopes for load and store operations, noting that color and depth/stencil are done by different stage/access pairs.

Debugging Tips

Note: There are more detailed recommendations below in [“Debugging Hazards”](#)

- Non-zero `Access info read_barrier` and `write_barrier` that do not include the current usage, likely reflect an incorrect destination mask (second scope).
- Zero (empty) `Access info read_barrier` and `write_barrier` likely reflect the absence of any barrier or an insufficient or incorrect source mask (first scope)
- Insert additional barriers with stage/access `VK_PIPELINE_STAGE_ALL_COMMANDS_BIT`, `VK_ACCESS_MEMORY_READ_BIT|VK_ACCESS_MEMORY_WRITE_BIT` for both `src*Mask` and `dst*Mask` fields to locate missing barriers.

- Use `vkSetDebugUtilsObjectNameEXT` to name created objects, as Vulkan handles are not invariant from run to run.

What is synchronization and why is it important?

Correct synchronization is needed to ensure correct results from Vulkan operations (whether graphical or computational). Modern graphics hardware is both parallel and pipelined, with various operations happening simultaneously for performance reasons. Different implementations may have differing pipeline architectures and/or memory hierarchies and caches. As such, apparently correct operation on a given Vulkan hardware or software implementation does not imply application correctness nor guarantee correct operation on a different implementation.

Vulkan defines a limited number of ordering guarantees, but for most operations, it is the application's responsibility to inform the implementation when ordering is required between operations. The need for such synchronization operations arises when the same region of memory is used by subsequent operations in different ways -- for example a mip-level being written by a blit operation, and then being used for sampled lookup by a shader.

If two uses are not guaranteed to operate sequentially, a data hazard exists. These hazards are:

RAW	Read-after-write	This occurs when a subsequent read operation uses the result of a previous write operation without waiting for the result to be completed.
WAR	Write-after-read	This occurs when a subsequent operation overwrites a memory location read by a previous operation before that operation is complete. (requires only execution dependency)
WAW	Write-after-write	This occurs when a subsequent operation writes to the same set of memory locations (in whole or in part) being written by a previous operation.
WRW	Write-racing-write	This occurs when unsynchronized subpasses/queues perform writes to the same set of memory locations.
RRW	Read-racing-write	This occurs when unsynchronized subpasses/queues perform read and write operations on the same set of memory locations

When *current* and *prior* are used to describe resource usage, this does not imply execution ordering on the Vulkan implementation. *Current* and *prior* refer to the sequence of Vulkan API entry calls in “Submission Order” and “Semaphore Order” (see the Vulkan Specification for a full description), as well as the order of Vulkan calls made by the API, assuming all “external synchronization” requirements are met.

When an application is insufficiently synchronized, data corruption of various types can occur. Blit, copy, or present operations may result in a destination image (or buffer) that is based on an incomplete source image (or buffer), or one partially updated from a previous state. Images that are used as both output attachments, and as shader inputs, may contain only partial results of output operations, or be overwritten while still by referenced as an input. Unsynchronized buffers may contain uninitialized or invalid constant, vertex, or index values corrupting in results of a draw or dispatch call. As with all parallel processing problems, these corruptions may appear only rarely, or sporadically, making them difficult to find and debug.

Correct synchronization is also important for application portability. Vulkan implementations may vary in the *effective* synchronization implicit in their software or hardware. This means that applications operating without correct synchronization may work correctly when tested on a given manufacture, model, or even driver version of a Vulkan implementation, but may fail with data corruptions (again, potentially rarely, and sporadically) on other versions or implementations. Even an optimization introduced by a new driver could expose a failure mode, indicating missing or insufficient set of synchronization operations.

Synchronization Operations

Synchronization operations create dependencies between operations accessing memory. These execution and memory dependencies are used to solve data hazards, i.e. to ensure that read and write operations occur in a well-defined order.

Write-after-read hazards can be solved with just an execution dependency, but read-after-write and write-after-write hazards need appropriate memory dependencies to be included between them. If an application does not include dependencies to solve these hazards, the results and execution orders of memory accesses are undefined.

While it is critical for data integrity that sufficient dependencies be defined to avoid these hazards, it is equally important that **excess** dependencies are not introduced, which would impact performance by overly serializing execution. Synchronization Validation can be used to test reductions in dependencies and verify that the less restrictive synchronization scheme is still correct.

Pointers to synchronization blogs/articles

[Hans-Kristian's in-depth blog post on Vulkan synchronization](#)

[Synchronization Examples](#)

[Video talk on "Keeping your GPU fed"](#)

[Understanding Vulkan Synchronization](#)

[Guide to Vulkan Synchronization Validation](#)

[Webinar: How to Validate Vulkan Synchronization \(Oct 2022\)](#)

[Synchronization2 Transition Guide](#)

Using Synchronization Validation

Before validating synchronization operations, resolve all validation errors from Standard Validation and Thread Safety. This will prevent wasted effort debugging hazards caused by invalid usage of the commands involved. It is possible to have entirely valid Vulkan command streams, and still have synchronization issues. However, invalid Vulkan commands may introduce synchronization issues by changing the effect of those commands.

Enabling synchronization validation.

Using Vulkan Configurator

The simplest way to enable Synchronization Validation is using [Vulkan Configurator](#). Select the “Synchronization Preset” to enable the recommended settings.

Using Layer Settings

This can also be done using your layer settings file, `vk_layer_settings.txt`

```
khronos_validation.enables =  
VK_VALIDATION_FEATURE_ENABLE_SYNCHRONIZATION_VALIDATION_EXT  
Khronos_validation.disable =  
VK_VALIDATION_FEATURE_DISABLE_OBJECT_LIFETIMES_EXT, VK_VALIDATION_FEATURE_DISABLE_API_  
PARAMETERS_EXT, VK_VALIDATION_FEATURE_DISABLE_CORE_CHECKS_EXT
```

Using Environment Variables

Synchronization validation can also be enabled through environment variables: (on non-Windows system replace “;” with “:”)

```
VK_LAYER_ENABLES=VK_VALIDATION_FEATURE_ENABLE_SYNCHRONIZATION_VALIDATION_EXT  
VK_LAYER_DISABLES=VK_VALIDATION_FEATURE_DISABLE_CORE_CHECKS_EXT;VK_VALIDATION_FEATURE_  
_DISABLE_OBJECT_LIFETIMES_EXT;VK_VALIDATION_FEATURE_DISABLE_API_PARAMETERS_EXT
```

The `VK_EXT_validation_features` extension can be used to enable Synchronization Validation programmatically at `CreateInstance` time, as shows in this code snippet:

```
VkValidationFeatureEnableEXT enables[] =  
{VK_VALIDATION_FEATURE_ENABLE_SYNCHRONIZATION_VALIDATION_EXT};  
VkValidationFeatureDisableEXT disables[3] = {  
    VK_VALIDATION_FEATURE_DISABLE_API_PARAMETERS_EXT,
```

```

    VK_VALIDATION_FEATURE_DISABLE_OBJECT_LIFETIMES_EXT,
    VK_VALIDATION_FEATURE_DISABLE_CORE_CHECKS_EXT
};
VkValidationFeaturesEXT features = {
    VK_STRUCTURE_TYPE_VALIDATION_FEATURES_EXT, nullptr, 1, enables, 3, disables
};
VkInstanceCreateInfo info = {};
info.pNext = &features;

```

Typical Usage

Given the challenge of debugging synchronization issues, applications should be tested within a debugger, set to break for any validation error. On Windows this can be done in Vulkan Configurator by selecting “Debug/Action” options “Break” and “Debug Output” causes Synchronization Validation to show the text of the error message in the debugger and pause program execution for each error. On all platforms, this can be accomplished by setting a debug breakpoint in a debug callback. The debug callback is defined using `vkCreateDebugUtilsMessengerEXT` with `VK_DEBUG_REPORT_ERROR_BIT_EXT` set in `VkDebugReportCallbackCreateInfoEXT::flags`. For each reported hazard, the provoking Vulkan call will be on the debugger call stack, aiding debugging.

Understanding Synchronization Validation Messages

Message Contents

All synchronization error messages begin with `SYNC-<hazard name>` where `<hazard name>` is one of the hazard types listed above.

The message body for each is constructed:

```
<cmd name>: Hazard <hazard name> <command specific details> Access info (<...>)
```

Command specific details typically include the specifics of the access within the current command.

Examples of typical command specific detail:

Command type	Details
Copy or blit	source or destination and region index
Draw or dispatch	Descriptor: binding, type Attachment: index and type Bound buffer: vertex or index
ImageBarriers	oldLayout, newLayout, subresource
Render pass	transitions oldLayout, newLayout load/store/resolve: attachment index, type, and operation
QueueSubmit	The command buffers and usage information for both the currently submitted command buffer and batch, as well as the previously submitted command buffer.

The `Access info` is common to all Synchronization Validation error messages. The fields in `Access info` are:

Field	Description
<code>usage</code> <code>executed_usage</code> <code>submitted_usage</code>	The stage/access using the memory range by the current command being recorded, invoked with <code>vkExecuteCommands</code> , or submitted to queue (respectively) (referred to as <code>usage</code> below)
<code>prior_usage</code>	The stage/access of the previous (hazarded) memory use (same naming as <code>usage</code>). The prior usage is defined as being earlier in the command buffer being recorded, earlier in Queue Submission Order, or by a previously called submission on a different queue.
<code>read_barrier</code>	For read <code>usage</code> , the list of stages with execution barriers between <code>prior_usage</code> and <code>usage</code>
<code>write_barrier</code>	For write <code>usage</code> , the list of stage/access (in <code>usage</code> format) with memory barriers between <code>prior_usage</code> and <code>usage</code>
<code>command</code>	The command that performed <code>usage</code> , <code>executed_usage</code> , <code>submitted_usage</code> , Or <code>prior_usage</code>
<code>command_buffer</code>	The command buffer in of a <code>executed_usage</code> , <code>submitted_usage</code> , Or <code>prior_usage</code>
<code>seq_no</code>	The one based index of <code>command</code> within the command buffer it is recorded within
<code>reset_no</code>	the reset count (one being the first record version) of the command buffer in which <code>command</code> is recorded
<code>queue</code>	The queue a prior command buffer was submitted to (Queue submit hazards only)

Recorded/Executed/Submitted Comparative Example

As noted, there are slight differences in the format of synchronization validation error messages. Below are three similar error messages.

For a hazard within a single command buffer, the `usage` references the stage/access which results from the current Vulkan API call. The `prior_usage` references the stage/access of a command recorded earlier in the current command buffer.

```
Validation Error: [ SYNC-HAZARD-WRITE-AFTER-READ ] Object 0: handle =
0xfa21a40000000003, type = VK_OBJECT_TYPE_BUFFER; | MessageID = 0x376bc9df |
vkCmdCopyBuffer(): Hazard WRITE_AFTER_READ for dstBuffer VkBuffer
0xfa21a40000000003[], region 0. Access info (usage: SYNC_COPY_TRANSFER_WRITE,
prior_usage: SYNC_COPY_TRANSFER_READ, read_barriers: VkPipelineStageFlags2(0),
command: vkCmdCopyBuffer, seq_no: 1, reset_no: 1).
```

Note that no command buffer information is given for either usage, as all memory operations involved are within the same, current command buffer. The `seq_no` also refers to the commands in the current version or the current command buffer. The `reset_no` is useful to determine where in application code flow this command buffer is being recorded.

A similar error (though a different hazard) shows a hazard between a secondary and primary command buffer. The `executed_usage` is the secondary command buffer with information given. The `prior_usage` reflects a command recorded earlier in the primary `command buffer`. If prior usage had been from an earlier `vkCmdExecuteCommands` call, a `command_buffer` field would have been included in the prior usage information.

```
Validation Error: [ SYNC-HAZARD-WRITE-AFTER-WRITE ] Object 0: handle =
0x21517094720, type = VK_OBJECT_TYPE_COMMAND_BUFFER; | MessageID = 0x5c0ec5d6
| vkCmdExecuteCommands(): Hazard WRITE_AFTER_WRITE for entry 0,
VkCommandBuffer 0x2150cfbee80[], Executed access info (executed_usage:
SYNC_COPY_TRANSFER_WRITE, command: vkCmdCopyBuffer, seq_no: 1, reset_no: 1).
Access info (prior_usage: SYNC_COPY_TRANSFER_WRITE, write_barriers: 0,
command: vkCmdCopyBuffer, seq_no: 1, reset_no: 11).
```

Finally, for a hazard reported during **queue** submission, the **submitted_usage** gives the stage access information for a command buffer in the current batch of the current **queue** submit call. The **prior_usage** reflects that stage/access of a command in an earlier command buffer (again in **Queue Submission Order**). A **command_buffer** field is provided for the prior usage. The “current” (actually the first use in the submitted command buffer is **provided** in the body of the error.

```
Validation Error: [ SYNC-HAZARD-WRITE-AFTER-READ ] Object 0: handle =
0x1febb508d20, type = VK_OBJECT_TYPE_QUEUE; | MessageID = 0x376bc9df |
vkQueueSubmit(): Hazard WRITE_AFTER_READ for entry 1, VkCommandBuffer
0x1febae67c50[], Submitted access info (submitted_usage:
SYNC_COPY_TRANSFER_WRITE, command: vkCmdCopyBuffer, seq_no: 1, reset_no: 2).
Access info (prior_usage: SYNC_COPY_TRANSFER_READ, read_barriers:
VkPipelineStageFlags2(0), queue: VkQueue 0x1febb508d20[], submit: 0, batch: 0,
batch_tag: 1, command: vkCmdCopyBuffer, command_buffer: VkCommandBuffer
0x1fec5015920[], seq_no: 1, reset_no: 2).
```

Understanding Synchronization Validation

When a Vulkan command that accesses memory is recorded or submitted for execution, the accesses for that command are tested for potential conflicts with previous accesses -- the hazards listed above. These accesses may be explicitly defined by command parameters, for example copy and blit source and destination. The accesses may be defined by the creation parameters of another Vulkan object (for example a `VkRenderPass`, `VkFramebuffer`, or `VkShaderModule`), or by descriptor or buffer binding. In addition to copy, blit, draw, and dispatch calls which clearly imply memory accesses, synchronization operations (barriers, events) and render pass operations (begin, next, end) may have implicit accesses. The user should not be surprised when hazards occur during (or due to) these other types of Vulkan commands.

When running Synchronization validation against your code, the current functionality will report hazards between accesses:

- Within the same command buffer
- Between primary and secondary command buffers
- Between command buffers submitted on the same or between queues
- Between operations recorded in submitted command buffers and swapchain operations

The error messages will list the command during which the hazard occurred, command specific details of which access within the command conflicts with a prior access, details about the current access, and the prior command and access with which the current

access conflicts. The prior command is identified by command name and the sequence number within the current command buffer, or within the previously submitted command buffer.

Stage/Access Usage Pairs

While the Vulkan specification typically lists pipeline stages and access types independently, only a small subset of the possible combinations of stage and access types are valid. (See *Table 4. Supported access types* in the [Access Types](#) section of the Vulkan API Specification). In error messages the stage/access pairs are given as `SYNC_<stage>_<access_type>` for brevity, denoting an access of `VK_ACCESS_<access_type>_BIT` occurring on stage `VK_PIPELINE_STAGE_<stage>_BIT`. (Note that for extension bits, the extension tag is appended to the `SYNC_...` string.). Additional `SYNC_...` strings are defined for implicit accesses without stage or access. For example `SYNC_IMAGE_LAYOUT_TRANSITION` denotes the read/write access implied by an image layout transition.

Hazard Detection

The Vulkan specification describes synchronization dependencies in terms of relationships between operations. Synchronization Validation however looks at the impact of synchronization operations on the safety (or correctness) of subsequent actions on resources, whether ranges of memory, or image subresource ranges and extents. For each Vulkan command that operates on memory, the prior state of affected memory ranges is inspected for the stage/access type of prior usage, the effect of synchronization operations on which subsequent usages are known safe relative to the prior accesses. Synchronization Validation error messages report hazards caused by the current Vulkan command's resource accesses relative to prior accesses for the same resources.

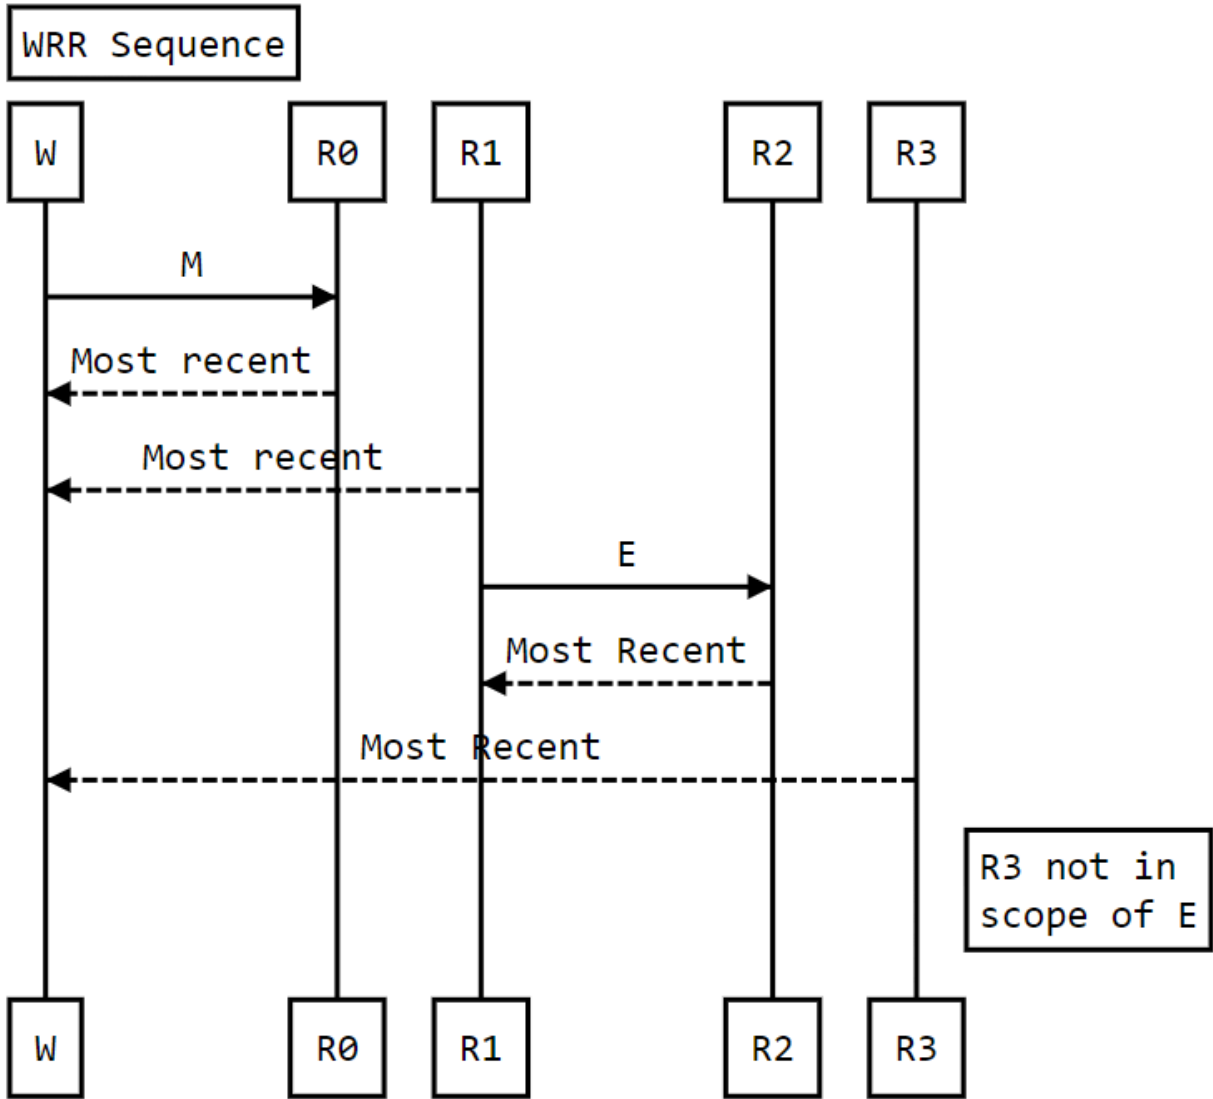
Most Recent Access

When reporting hazards, Synchronization Validation it is sufficient to only inspect the *most recent* access for a given memory or image subresource range. Since, all prior hazards are assumed to have been reported, tracking state prior to the most recent access is unneeded. Additionally, when reporting hazards with recorded command buffers at `vkExecuteCommands` or `vkQueueSubmit` time, only the *first* access for a given memory location need be checked.

Read Operations

For read operations the most recent access rules apply to prior reads with execution barriers (or ordering) relative to the current read. The prior write access is only considered the most recent access if no intervening prior read has occurred that *happens-before* the current read. Consider the following sequence of access and barriers (listed in submission order) acting on the same memory address:

Operation	Description
W	write operation
M	memory barrier guarding access at R0
R0	first read operation
R1	second read operation
E	execution barrier such that R2 <i>happens-after</i> R1
R2	third read operation
R3	fourth read operation with stage not in second execution scope of E

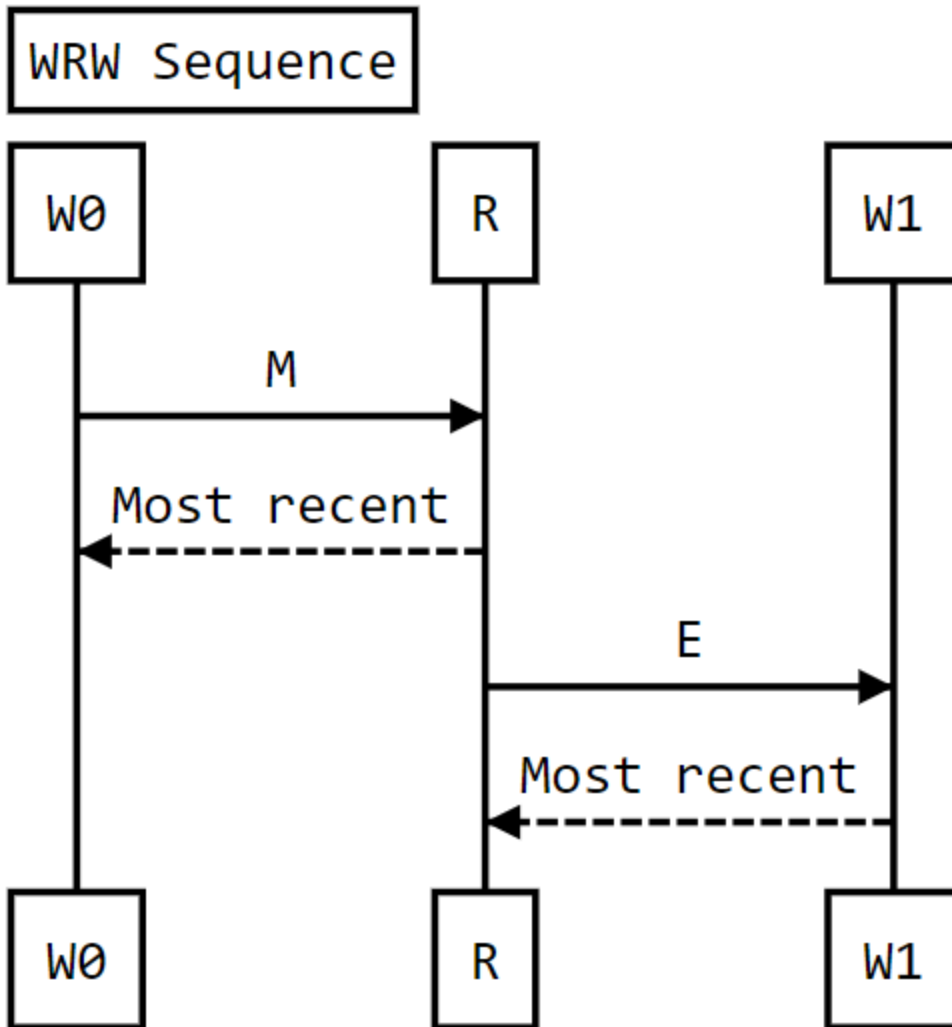


Write Operations

For write hazard checks in a given range of memory addresses, if there are intervening read operations between the current write and the most recent previous write, these intervening read operations are considered the most recent access. In that case, write-after-write checks are not done.

Consider the following sequence of operations on the same memory address:

Operation	Description
W0	first write operation
M	memory barrier guarding access at R
R	read access
E	execution barrier guarding access at W1
W1	second write operation



In this case, a read-after-write check is done for R based on W_0 and M , and a write-after-read check is performed on W_1 based on R and E . W_1 is not checked against W_0 for write-after-write. If W_0, M, R is not a hazard, this guarantees W_0 is available and visible to R , and thus to any operations that *happen-after*. As such, the correctness of R, E, W_1 depends solely on those operations. The correctness of the entire sequence can be assured by pairwise hazard checks.

Debugging Hazards

While full details of the current resource access are available for the current Vulkan command, the information for locating `prior_usage` are more limited. Most applications (apart from debug and replay tools) will not have the sequence number (and traceback) of commands within the current command buffer, so the location of the access w.r.t. the `prior_usage` may take some effort to determine. Additionally, the `command` correlated with the `prior_usage` will likely be correct, and indicate that the synchronization operations between `command` and the current command are missing, incorrect, or incomplete.

An important note (repeated here from the quickstart), is that for *memory* barriers stage masks are *not* logically extended to logically earlier and later stages, but only apply to the specific stages specified. *Execution* barriers with a `dstStageMask` of `VK_PIPELINE_STAGE_VERTEX_SHADER_BIT` do imply an execution barrier including all later stages (for example `VK_PIPELINE_STAGE_FRAGMENT_SHADER_BIT`). However, assuming an access specified by `dstAccessMask` of `VK_ACCESS_SHADER_READ_BIT` would only define a memory barrier with the *vertex shader stage/shader read stage/access*. No *memory barrier* would be established with the *fragment shader stage* for the specified access.

For developers intimate with the details of their rendering pipeline, this information may be sufficient to directly identify the cause of the hazard. However, the following approaches may help refine the cause and fix for identified hazards.

Debugging Using Access info information

The access information gives several useful pieces of information to help locate the missing synchronization operation. Using this information to find the synchronization can depend on the type of hazard and `prior_usage`.

Hazards from Missing or Incomplete Barriers

Read and write barriers in the error message can help identify the synchronization operation (either subpass dependency or pipeline barrier) with insufficient or incorrect destination stage/access masks (second scope). Insufficient barriers may be present due to several reasons which may be identified by examination.

Zero (empty) Read and Write Barriers

This can indicate that either no barrier with correct scopes was applied to the reported resource, or applied to the wrong resource, or subresource range. This can be found by searching for a synchronization operation which includes the needed stage(s)/access(es) to avoid the reported hazard, and noting the resources affected.

Alternatively, this can indicate that any barrier applied to the resource had an source stage/access mask set (first scope) that did not include the prior usage with which the current, executed, or submitted usage hazards. This might be identified by identifying an earlier synchronization operation which includes destination stage/access masks (second scope) which include the current (executed, or submitted) usage, with insufficient first scope to include the prior usage.

Non-Zero Barriers

Access info `read_barrier` and `write_barrier` values that are non-zero but do not include the stage/access of the current usage, likely reflect an incorrect destination mask (second scope) in a prior synchronization operation. This may be found by identifying the earlier synchronization operation matching the `read_barrier` and `write_barrier` values reported.

Hazards vs. Prior Image Layout Transitions

If the `prior_usage` is `SYNC_IMAGE_LAYOUT_TRANSITION`, the `write_barrier` should help identify the `vkCmdPipelineBarrier` or `VkSubpassDependency`, of the `prior_usage`. The stages/access listed correlating to the `dstStageMask` and `dstAccessMask` for the prior command should match the `write_barrier` value. The fix is likely (barring missing dependency chain `dstStageMask` and chaining `srcStageMask` bits) to ensure that the current usage is added to the `dstStageMask` and `dstAccessMask` fields of the barrier or renderpass operation that performed the layout transition.

Hazards at Image Layout Transitions

If the hazard occurs at a layout transition, you likely have no further to look than the `srcStageMask` and `srcAccessMask` of the current operation to find the missing or incorrect stage or access needed to guard the layout transition. Depending on code structure, the optimal solution may be to provide barriers at an earlier stage (when the presence or absence of a potential hazard would be known).

Hazards between buffer and/or image resource uses

When a resource changes between roles (being written or read) without an appropriate synchronization operation, a hazard will occur. Common situations are transitions between transfer and rendering (or compute) operations. If the `write_barrier` or `read_barrier` fields are non-zero, the actual error is likely a malformed barrier between usages. The `command` and `prior_usage` fields should help identify the previous access and the barrier field can be used to find the existing barrier, event, or renderpass operation similarly to finding prior image layout transitions. Barriers constructed from a chain of dependencies may be more difficult to back trace in this way and may require a different approach.

Method of Bisection Using Additional Barriers

In addition to back-tracking from `Access info` to find the missing or incomplete barriers, one can debug these issues by a simple bisection search. To determine where in program flow a missing barrier occurs, simply insert a serializing debug barrier (described below) within your application *prior* to the hazard. If the hazard is removed, then the cause of the hazard is prior to the location of the debug barrier. If the hazard remains, the cause of the hazard is subsequent to the debug barrier.

A debug barrier is either a `VkCmdPipelineBarrier` or `VkSubpassDependency`, that specifies that all execution and access after the barrier *happen-after* all execution and access prior to it. In both cases set `srcStageMask` and `dstStageMask` set to `VK_PIPELINE_STAGE_ALL_COMMANDS_BIT`. For the pipeline barrier one `VkMemoryBarrier` is required with both `srcAccessMask` and `dstAccessMask` set to `VK_ACCESS_MEMORY_READ_BIT|VK_ACCESS_MEMORY_WRITE_BIT`. The corresponding fields within `VkSubpassDependency` are set to the same value. If a `VkCmdPipelineBarrier` debug barrier is to be used during a subpass, it must all be specified as a self-dependency for that subpass. (See [Subpass Self-dependency](#) in the Vulkan API Specification)

Note: this is for debug only and should not be left in production code as this could have significant performance impact as the barrier serializes all GPU operations before and after the inserted barrier.

Identifying Affected Resources and Operations

One important challenge is to work from the Vulkan object handle and command information (for example `seq_no: 1` and `reset_no:`) to the application code path that generated the command. There are two linked problems identifying the application resource or data structure associated with a given Vulkan resource, and tracking operations on that resource.

Getting Consistent Resource Identification

Within the Vulkan API, created resources are not guaranteed to be invariant from run to run of an application. This means that the handles in synchronization error messages will not be consistent during the debugging process. One solution is to use `vkSetDebugUtilsObjectNameEXT` and `vkSetDebugUtilsObjectTagEXT` to name or tag all Vulkan resources (buffer, images, command buffers, etc.) involved in the reported Hazard. This requires application code modification to generate and attach this tagging information. Naming information is then reflected in the Hazard messages, and may provide enough information to suggest useful “By Code Inspection” techniques (below) to pursue.

Tracking Operations For a Given Resource

Once a named handle is identified, the operations on that handle must be tracked. To do this the handle (and/or any application data structure) needs to be identified at the time the name (or tag) is assigned. Given that the `prior_usage` field `command` is known, breakpoints can be set for calls to that Vulkan entrypoint which reference the handle matching the name or tag of interest. Noting that usage of a resource may be indirect (for example through a descriptor), additional tracking of command referencing that indirection may be required (including additional naming and breakpoints).

Using Code Inspection

Sometimes the cause of a synchronization error can be found by simple inspection of the application or engine code.

Look near the stack trace location

It's possible that an incorrectly specified barrier is literally in or near the debugger stack trace. Look for the type of issues with stage and access masks noted above errors (e.g. missing specific stages for a given usage memory access).

Identifying Incomplete Existing Barriers

If not near the current code location, it is possible to identify an existing barrier that is incomplete. To do this, search the code for references to `VK_PIPELINE_STAGE_*` or `VK_ACCESS_*` flags that match either the *current* usage (check `dst*Mask`) or *prior* usage (check `src*Mask` fields) but that do not include the correct flags for the opposite usage. For example, there may be a synchronization operation with second (destination) scope stage/access flags that match the current usage, but the first (source) scope stage/access flags not matching prior usage and vice versa. Always remember to check that the code didn't incorrectly assume logical stage expansion applied to memory access barriers.

Examining Resource Use Transitions

A common element to many Vulkan applications or engines is a subsystem that tracks the current use of a given resource. For example is the resource a transfer destination, an output attachment, or a shader input. Examining the application's or engine's view of the resource of the Hazard, one can note if the application either 1) hasn't transitioned the intended use of the resource, 2) has missing or incorrect synchronization, or the output of a pipeline. Note the usage/ownership the application internal data structures believe the resource to be in. Does that match the usage/API functional element making the current operation? Look for the code that does usage/ownership transition for application resources for the type of resource (and usage) involved. Check the synchronization operations that would have been applied for the resource usage/ownership changes reflected in the current usage and prior usage provided by the Hazard error message

Optimizing Synchronization with Synchronization Validation

While synchronization is required to avoid data corruption within a Vulkan application, excessive synchronization can negatively impact performance by introducing stalls in command execution. Certainly any synchronization operations using `VK_PIPELINE_STAGE_ALL_COMMANDS_BIT`, `VK_ACCESS_MEMORY_READ_BIT`, or `VK_ACCESS_MEMORY_WRITE_BIT` masks, or device or queue wait operations should be targeted for elimination. Synchronization Validation can be used to identify the minimum needed barriers and dependencies. To do this, one can reduce the scopes of barriers and dependencies and note the hazards that are reported by Synchronization Validation. By adding just the barriers needed to eliminate the hazard error message, you can establish a minimally synchronized set of barriers and dependencies. Clearly one must be careful to exercise all possible Vulkan command sequences capable of producing different hazards, to assure that the narrowed set maintains correctness.

The minimal set of synchronization operations established may be quite broad (in terms of a large number of stages and accesses guarded). This can occur when the Vulkan usage pattern is highly variable before and after the synchronization operations. In these cases, it may be useful for the application to customize the barriers for the various use patterns. In addition, reordering processing steps to reduce the need for synchronization, may be of value.

In assessing the need to further optimize synchronization, performance benchmarks, and GPU usage tools should can and should be used. At each step however, testing with Synchronization Validation can assure that correctness is maintained.

For Further Information and Current Status

Please contact LunarG to let us know what you think of Synchronization Validation or to offer suggestions for future releases. The best way to reach us is via our GitHub repository using the link shown below:

<https://github.com/KhronosGroup/Vulkan-ValidationLayers>

Status, known limitations, and ongoing work on this project can be followed on the Synchronization Validation project page:

<https://github.com/KhronosGroup/Vulkan-ValidationLayers/projects/5>

Revision history

Revision Date	SDK Release	Comment
August 28, 2020	SDK 1.2.148.1	Alpha version, initial release date
January 11, 2021	SDK 1.2.162.1	V1.0: Supporting single command buffer hazard detection. Remove alpha specific comments, updated enable strings to standard versions, added project information
Jan 18, 2024	SDK 1.3.275.0	V2.0: Updated with additional functional and debugging information including Queue Submit functionality.
Feb 1, 2024	SDK 1.3.275.0	V2.1: Corrected errata regarding seq_no and reset_no basis, and recommended configuration