



# VMware vSphere 7.0 U2 Tagging Best Practices

Performance Study - March 1, 2023



VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 [www.vmware.com](http://www.vmware.com)  
Copyright © 2023 VMware, Inc. All rights reserved. VMware and the VMware logo are registered trademarks or trademarks of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMware products are covered by one or more patents listed at [www.vmware.com/go/patents](http://www.vmware.com/go/patents).

## Table of Contents

1. Executive Summary.....	4
2. Introduction.....	4
3. Terminology and Scope of Paper.....	5
4. Tagging APIs: Java.....	6
4.1. Create tag (assuming the category is already created).....	6
4.2. Associate tag(s) with VM(s).....	8
4.2.1. Single tag to single VM.....	8
4.2.2. Single tag to multiple VMs.....	9
4.2.3. Multiple tags to a single VM.....	11
4.2.4. Multiple tags to multiple VMs.....	12
4.3. Associate tags with other objects.....	18
4.3.1. Query VMs associated with tags.....	18
5. Tagging APIs: PowerShell.....	23
6. Linked Mode.....	25
6.1. Replication Performance.....	25
6.2. Comparison to standalone vCenter tagging performance.....	26
7. Tags vs. Custom Attributes.....	27
8. Scale Numbers for Good Performance.....	28
9. Concluding Remarks.....	29
10. References.....	31
11. Appendix.....	33
11.1. Experimental setup and software versions.....	33
11.2. Java code examples.....	33
11.2.1. Example J0: VM DynamicID and Tag TagID.....	33
11.2.2. Example J1: Retrieving category/tag IDs from category/tag names.....	34

11.2.3. Example J2: Create Category .....	35
11.2.4. Example J3: Create Tag .....	35
11.2.5. Example J4: Associating a tag with a VM (attach()) .....	36
11.2.6. Example J5: Associating a tag with multiple VMs (attachTagToMultipleObjects()) .....	36
11.2.7. Example J6: Associating multiple tags with multiple VMs (attachMultipleTagsToObject()) .....	37
11.2.8. Example J7: List tags associated with a VM .....	37
11.2.9. Example J8: List VMs associated with a tag.....	37
11.2.10. Example J9: List tags associated with a group of VMs .....	38
11.2.11. Example J10: List VMs associated with a group of tags.....	38
11.2.12. Example J11: Retrieving tag information (e.g., name, description) from the tag ID.....	38
11.2.13. Example J12: Retrieving all associations using the pagination API.....	39
11.3. PowerShell code examples .....	40
11.3.1. Example P0: Creating VMID Objects and retrieving Tag IDs .....	40
11.3.2. Example P1: List tags associated with a VM .....	41
11.3.3. Example P2: List VMs associated with a tag .....	41
11.3.4. Example P3: List tags associated with a group of VMs .....	41
11.3.5. Example P4: List VMs associated with a group of tags .....	41
11.3.6. Example P5: Retrieve Associations using Pagination API .....	42

# 1. Executive Summary

Writing code to use VMware vSphere® tags can be challenging in large-scale environments. In this technical paper, we discuss the scalability limits of tags and give some tips and tricks for writing performant tagging code in Java or PowerShell.

This paper is an updated version of the tagging whitepaper for the vSphere 7.0 U2 release. We describe some performance improvements (relative to 7.0 U1) in tag association APIs and show data regarding Enhanced Linked Mode performance. The limits for the number of tag associations vSphere can support are the same as for 7.0 U1, although the performance for some APIs is improved.

## 2. Introduction

vSphere 5.1 introduced an inventory tagging feature available in all later versions of vSphere, including vSphere 7.0 U2. Tags let datacenter administrators organize different vSphere objects like datastores, virtual machines, hosts, and so on. This makes it easier to sort and search for objects that share a tag, among other things.

In vSphere, tags are organized into categories. A category is typically used for a high-level description. For example, a category might be “OS type” or “application name.” Within a category, the tags are the distinct values for that category. For example, if the category is “OS type,” then the tags might be “Linux” or “Windows 2016.” If the category is “application name,” the tags might be “DB,” “Middleware,” or “vCenter Server.” If the category is “CPU type,” the tag values might be “AMD,” “Intel Broadwell,” or “Intel Cascade Lake.”

Writing code to use tags can be challenging in large-scale environments. In this paper, we discuss the scalability limits of tags and give tips and tricks for writing performant tagging code in Java or PowerShell. While the performance results in this paper are specific to the hardware configuration and vCenter Server Appliance (VCSA) configuration in our lab, the fundamental trends will be the same in customer environments.

We first discuss the scalability limits for tags in vSphere 7.0 U2. We then provide best practices for writing code for tagging.

### 3. Terminology and Scope of Paper

Before describing the scale limits for tagging, we first define some terms:

- A *category definition* is a grouping of tags, for example, “OS type” or “application name.”
- Given the previous category names, example *tag definitions* might be “Windows” for the category “OS type” or “SQL Server” for the category “application name.”
- A *tag association* is the mapping between a tag and an object (like a virtual machine or datastore). For example, if tag “Windows” is applied to VM “vm-111,” then that is one association, and if tag “Windows” and tag “SQL Server” are both applied to “vm-111,” then that is two associations. If a user associated 25 tags (“Windows,” “SQL Server,” “Pasadena Datacenter,” and so on) with 1,000 VMs, that would be 25,000 tag associations.

There are no *hard* limits to the number of tags, categories, and tag associations. However, increasing the number of tags, categories, or tag associations can impact performance. In this paper, we have tried to provide information you can use to understand the performance limits in your environment. We give best practices for how to write scripts to create tag associations and retrieve association information from vCenter. We focus on Java- and PowerCLI-based programs, although the same principles apply to other languages like Python. Moreover, we focus primarily on three operations:

1. Creating tag definitions
2. Associating (attaching) tags to entities like VMs or datastores
3. Retrieving tag associations

These three operations are the most common tasks performed by scripts or plugins.

## 4. Tagging APIs: Java

In vSphere, a tag can be created only as part of an existing category. Thus, a category must be created before a tag can be added to it. The performance considerations for creating categories are like those for creating tags, so we discuss only tag creation in this section. For a simple example of creating a category using Java, please see 11.2.3. Example J2: Create Category in 11. Appendix.

Also, see the [API documentation for vCenter management/vSphere 7.0 U2](#) and the [vSphere Automation SDK Programming Guide 7.0 Update 2](#).

### 4.1. Create tag (assuming the category is already created)

Sample code for creating a category and creating a tag within that category is given in the 11. Appendix (please see 11.2.3. Example J2: Create Category and 11.2.4. Example J3: Create Tag).

Tag creation requires a name and a description. The list of associable types (that is, what types of objects, like VMs or hosts, can be associated with this tag) is optional: if it is not supplied, then a tag can be associated with any object. The category for a tag must specify the cardinality of the category. A description of cardinality is given in the performance team [tagging blog \[1\]](#), but briefly, *multiple cardinality* means that multiple tags (say, “Alice” and “Bob”) from a given category (say, “Owners”) can be assigned to a given object (like a VM). If the cardinality is one, then only one tag within a category may be applied to a given entity. For example, if the category is “OS,” then only one tag from that category (say, “Ubuntu Linux” or “Windows Server 2016”) may be applied to the VM.

The tag creation API allows a user to specify a single tag to be created—in this API, only one tag can be created at a time. The time to create a tag varies with the number of tags per category. For example, creating many tags under a single category is slower than creating a smaller number of tags in multiple categories.

Figure 1 shows the latency of creating 8,000 tag definitions under a single category. For example, if the category name were “Application,” the tag definitions might be “SQL Server,” “SAP Hana,” “Oracle,” “Kafka Messaging,” and so on. Creating 8,000 tags under 1 category increases slightly from the first tag to the last tag. While the times do vary, as seen in the figure, the average latency to create a tag is roughly 10 milliseconds. This result is in line with the results from 7.0 U1 and represents a performance improvement of over 6.7, which showed a much more dramatic linear increase in latency as the number of tags was increased. This improvement is due partly to caching tag information and database query optimizations.




Figure 1: Creating 8,000 tags under 1 category. In 7.0 U2, the time to create a tag stays roughly constant with the number of tags – on average, about 10 milliseconds.

Figure 2 shows the latency of creating 8,000 tag definitions under 200 categories. In each case, there are 40 tag definitions per category. As with the previous figure, the latencies show some variation, but the average latency to create tags is roughly constant: the time to create 40 tags per category is around 8 milliseconds per tag. This number is on par with creating all tags within a single category. As we had described in our [6.7 whitepaper \[4\]](#) one option to speed up tag creation is to use multiple clients to create tags concurrently. In our labs, we have tested up to 8 clients with promising speedups: up to 7x for 8 clients [5]. Although performance in 7.0 U2 is similar with one category vs. many categories, multiple categories allow for better parallelism when using multiple clients. Thus, in 7.0 U2, as with 6.7, if practical, we recommend multiple categories with fewer tags per category vs. a small number of categories with many tags in them.

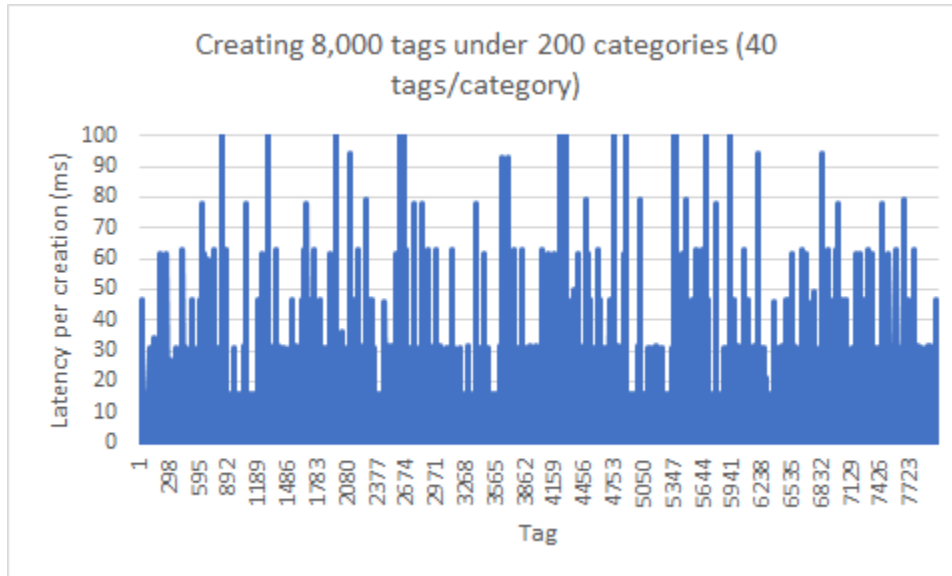


Figure 2: Creating 8,000 tags under 200 categories. Spreading tags among categories speeds up tag creation. This makes it roughly constant for small numbers of tags per category.

## 4.2. Associate tag(s) with VM(s)

There are 3 ways to attach a tag to a VM:

1. `attach(tagID, VM_ID)`
2. `attachTagToMultipleObjects(tagID, list of VM_IDs)`
3. `attachMultipleTagsToObject(list of tagIDs, VMID)`

Each of these calls is useful in different scenarios. The main difference is how many calls to the tagging service are required.

### 4.2.1. Single tag to single VM

When attaching a single tag to a single VM, it is easiest to use a simple `attach()` call (option 1 above). Please see 11.2.5. Example J4: Associating a tag with a VM (`attach()`) in 11. Appendix where we assign a tag “Windows 10” from the category “GuestOS” to a VM. We recommend keeping a map of the tag ID to the tag name and the VM ID to the VM name, if possible, to avoid calling vCenter or the tagging service each time such a mapping is needed. Pseudocode for using `attach()` is given here:

1. Get the tag ID (`tagID`) from tag name.
2. Get the VM ID (`VM_ID`) from the VM name.
3. `attach(tagID, VM_ID)`.

Attaching a single tag to a single VM typically takes around 18 milliseconds, less than the 25 milliseconds average in 7.0 U1. Moreover, unlike in 6.7, in which the time varies as the number of tag associations increases, with 7.0 U1 and 7.0 U2, the time is roughly constant.



Figure 3 shows the latency to attach a tag to a VM as we increase the number of VMs from 1 to 10,000. The latency varies, but most `attach()` calls take around 18 milliseconds. The total end-to-end time is 3.1 minutes. The `attach()` call is faster in 7.0 U2 vs. 7.0 U1 due to caching privileges.




Figure 3: Using `attach()` to associate 1 tag with each of the 10,000 VMs. The time for an individual `attach()` is roughly constant even as we increase the number of VMs.

#### 4.2.2. Single tag to multiple VMs

When attaching a single tag to multiple VMs, it is easiest to use option 2, `attachTagToMultipleObjects()`. This call amortizes the cost of attaching the tag over multiple VMs. For attaching the same tag to 10 VMs (VMID\_0-VMID\_9), only one call to the tagging service is required, as opposed to 10 if we had used an `attach(tagID, VMID_N)` call in a loop over each VM. See 11.2.6. Example J5: Associating a tag with multiple VMs (`attachTagToMultipleObjects()`) in 11. Appendix for an example of attaching a single tag to multiple VMs using `attachTagToMultipleObjects()`.

Pseudocode is given here:

1. Get tag ID (String tagID)
2. Get list of VM Dynamic IDs (List<DynamicID> VM\_IDs)
3. `attachTagToMultipleObjects(tagID,VM_IDs)`

Attaching a single tag to multiple VMs using `attachTagToMultipleObjects()` is more efficient than looping over every VM and calling `attach()`. Table 1 compares the latency of attaching a single tag to multiple VMs using `attachTagToMultipleObjects()` vs. using `attach()`. As noted in the caveat above, in the case of `attachTagToMultipleObjects()`, we attach the tag to 2,000 VMs simultaneously. In the case of `attach()`, we sequentially attach the tag to each VM. The performance for `attach()` is slightly worse than in 6.7 U1 due to caching introduced after an association is created, but the creation and list APIs are faster overall, as we indicate in other sections of this paper. In 6.7, a 5 millisecond sleep is required between iterations to avoid hitting a throttle within the vapi-endpoint process. In 7.0 U2 (as was the case in 7.0 U1), no such sleep is required, so with an appropriate parallelism of client-based calls, we can match or exceed the throughput of 6.7. In 7.0 U2, `attachTagToMultipleObjects()` is substantially faster (> 120%) than both 6.7 U1 and 7.0 U1 as a result of caching permissions information across tagging calls.

Operation	Number of iterations	Latency for 1 tag, 20,000 VMs
<code>attach()</code>	20,000 (1 per VM)	452 seconds
<code>attachTagToMultipleObjects()</code>	10 (2,000 VMs at a time)	109 seconds*

Table 1: `attach()` vs. `attachTagToMultipleObjects()`

\* The `attach()` call is faster in 7.0 U2 (109 seconds) vs. 7.0 U1 (265 seconds) due to caching privileges.

#### 4.2.2.1. Chunking VMs in calls to `attachTagToMultipleObjects()`

We suggest putting no more than 2,000 VMs per `attachTagToMultipleObjects()`. Exceeding this number may cause errors, since the tagging service has some limits regarding the number of items in a request. To see if this issue is occurring in your system, look for this kind of error in the log file for the vapi-endpoint service (that is, `/var/log/vmware/vapi/endpoint/endpoint.log`):

```
Received request (910,429 bytes) is bigger than the allowed request size (204,800 bytes) - request blocked!
```

### 4.2.3. Multiple tags to a single VM

When attaching multiple tags to a single VM, it is best to use option 3, `attachMultipleTagsToObject()`. We can attach 10 tags (tag\_0-tag\_9) to a single VM with one tagging service call instead of the 10 that would be required if we used `attach(tag_N,VMID)` in a loop over each tag. We give an example of attaching multiple tags to VMs in section 11.2.7. Example J6: Associating multiple tags with multiple VMs (`attachMultipleTagsToObject()`).

Here is the pseudocode to accomplish this:

1. Get tag IDs and put them into a list (`List<String> tagIDs`)
2. Get VM ID (`VM_ID`)
3. `attachMultipleTagsToObject(VM_ID,tagIDs)`

We recommend using no more than 2,000 tags in a single call for the best performance.

Figure 4 shows the latency of attaching 1 through 200 tags to a single VM using `attachMultipleTagsToObject()`. The latency is nearly constant until around 25 tags, and then it starts to grow linearly. There is some run-to-run variability, which is why the result with 45 tags is slightly worse than with 50 tags. After 150 tags, the incremental overhead is linear but with a slight slope. These latencies are on par with 7.0 U1 and substantially improved relative to 6.7.




Figure 4: `attachMultipleTagsToObject()` for 1-200 tags and 1 VM. The latency is fairly steady until around 50 tags when it starts to grow linearly.

Figure 5 compares the latency of attaching multiple tags (1 up to 2,000) to a single VM using `attachMultipleTagsToObject()` vs. using serial `attach()` calls vs. `attachTagToMultipleObjects()`. As the figure shows, each method works well with a small number of tags, but `attachMultipleTagsToObject()` works best as the number of tags grows. For a small number of tags (< 50), the latency of `attachMultipleTagsToObject()` is around 200 milliseconds. For 2,000 tags, the latency is around 500 milliseconds. The `attach()` and `attachTagToMultipleObjects()` calls are better than 7.0 U1 and slightly slower per-iteration vs. 6.7, as described earlier. When attaching 2,000 tags to 1 VM, both `attach()` and `attachTagToMultipleObjects()` must be called 2,000 times, paying this extra cost on each iteration. In contrast, `attachMultipleTagsToObject()` is significantly faster in 7.0 U2 vs. 6.7, primarily due to caching tag information in vCenter, and it is the preferred API for this use case.




Figure 5: Attaching 1 through 2,000 tags to 1 VM. `attachMultipleTagsToObject()` is well suited for this use case.

#### 4.2.4. Multiple tags to multiple VMs

The previous cases considered assigning one tag to multiple VMs or multiple tags to one VM. In this section, we discuss assigning multiple tags to multiple VMs.

#### 4.2.4.1. 10 tags associated with 1,000 VMs (10,000 tag associations)

Suppose we had a group of 10 tags (for example, “Windows 2019 Server,” “DB app,” etc.), and we wanted to attach them to a group of 1,000 VMs (for example, VM\_0, VM\_1, ..., VM\_999). As the preceding discussion suggests, there are two ways to accomplish this:

1. `attachTagToMultipleObjects()` Pseudocode

```
Get tag IDs → List<String> tagIDs
Get VM IDs → List<DynamicId> VM_IDs
for (String tagID : tagIDs) {
    attachTagToMultipleObjects(tagID, VM_IDs)
}
```

2. `attachMultipleTagsToObject()` Pseudocode

```
Get tag IDs → List<String> tagIDs
Get VM IDs → List<DynamicId> VM_IDs
for (DynamicID VM_ID : VM_IDs) {
    attachMultipleTagsToObject(VM_ID, tagIDs)
}
```

We give more complete code snippets in the 11. Appendix.

In table 2 below, we show the relative performance of these two calls for assigning 10 tags to 1,000 VMs from our internal testing labs.

API Call	Number of iterations	End-to-end latency	Time per iteration
<code>attachTagToMultipleObjects()</code>	10	53 seconds	5 seconds
<code>attachMultipleTagsToObject()</code>	1,000	16 seconds	16 milliseconds

Table 2: Latency for assigning 10 tags to 1,000 VMs (a total of 10,000 associations). The end-to-end latency depends on the per-iteration cost.

As the table indicates, for this modest number of tags and VMs, `attachMultipleTagsToObject()` is faster than `attachTagToMultipleObjects()`, even though `attachMultipleTagsToObject()` executes more iterations. In the first case, each iteration is creating 1,000 tag associations (that is, 1 tag is being associated with 1,000 VMs). In the second case, each iteration is creating 10 associations (that is, associating 10 tags to 1 VM). Once again, the 7.0 U2 results are better than 7.0 U1 due to the caching of privileges.

#### 4.2.4.2. 15 tags associated with 5,000 VMs (75,000 tag associations)

In the previous example, the number of associations (10,000) was small enough that the latency to create each association was constant. However, the latency to attach tags to VMs increases as the number of associations increases.

##### 4.2.4.2.1. *attachTagToMultipleObjects()*

Figure 6 shows the latency of attaching multiple (1, 5, 10, 15) tags to up to 5,000 VMs using `attachTagToMultipleObjects()`. This creates up to 75,000 tag associations. To conform to the limit of 2,000 objects per call, for this example, we divided the 5,000 VMs into 2 batches of 2,000 VMs and 1 batch of 1,000 VMs. We then looped through the tags and attached the tag to each batch. Here is some pseudocode to illustrate this:

```
// (not shown): Get tag IDs into List<String> tagIDs
// (not shown): Get VM IDs into List<DynamicId> VM_IDs
int chunkSize = 2000
int numChunks = ((VM_IDs.size()/2000) + 1)
for (String tagID : tagIDs) {
    for (int j = 0; j < numChunks; j++) {
        int index = j * chunkSize
        // Make sure we don't exceed the bounds of the VM ID list
        end = max(VM_IDs.size(),index+chunkSize);
        attachTagToMultipleObjects(tag, VM_IDs.sublist(index, end));
    }
}
```

In the graph below, the time taken to attach the first tag to all VMs is around 26 seconds, and the time to attach the 15<sup>th</sup> tag to all VMs is also around 26 seconds. This is significantly faster than both 6.7 and 7.0 U1 (7x better than 6.7 and 3x better than 7.0 U1). The end-to-end time to attach all 15 tags is the sum of the time for each iteration, around 400 seconds compared to 3,000 seconds for 6.7 and 1,000 seconds for 7.0 U1. This experiment creates 75,000 associations. Results of experiments like this are one reason we feel comfortable suggesting higher limits for the number of supported tag associations in 7.0 U1 and 7.0 U2 (see 8. Scale Numbers for Good Performance). In general, we can exceed the out-of-the-box limits of the API by using batches of VMs, though the higher above the limit we go, the slower the API and UI will perform.

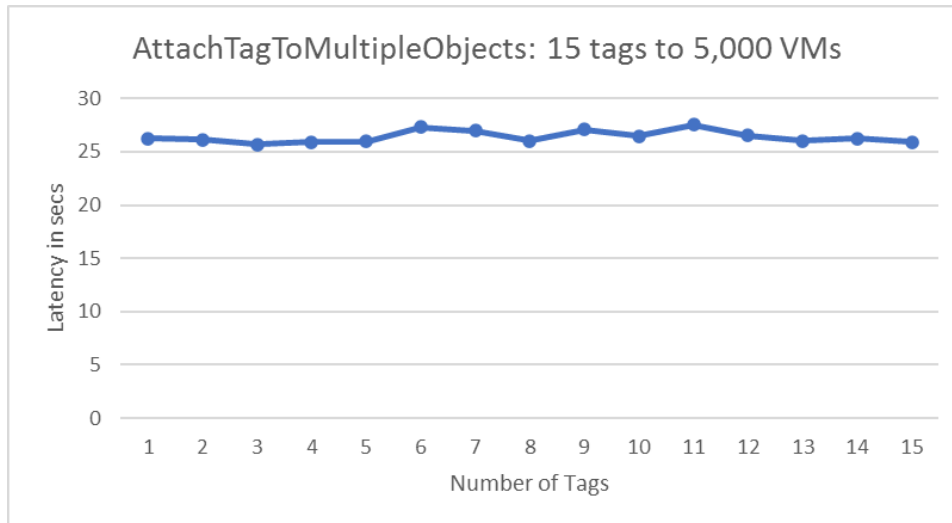


Figure 6: attachTagToMultipleObjects(): 15 tags to 5,000 VMs.

#### 4.2.4.2.2. attachMultipleTagsToObject()

We can attach 15 tags to 5,000 VMs using `attachMultipleTagsToObject()` instead. The results are shown in figure 7. To use `attachMultipleTagsToObject()`, we loop over VMs one at a time and attach 15 tags at once:

```
// (not shown) Get tagIDs into List<String> tagIDs
// (not shown) Get VM IDs into List<DynamicId> VM_IDs
for (DynamicID VM_ID : VM_IDs) {
    attachMultipleTagsToObject(VM_ID, tagIDs)
}
```

As figure 7 indicates, the time to attach 15 tags to the first few VMs is typically between 12 and 30 milliseconds, and the time to attach 15 tags to the 5,000<sup>th</sup> VM is also between 12 and 30 milliseconds. The total time to attach 15 tags to 5,000 VMs is the sum of the latency of each iteration, around 81 seconds, which is over 20% faster than for 7.0 U1.




Figure 7: `attachMultipleTagsToObject()`. This call is roughly constant with the number of objects.

**Caveat about the number of objects or tags in a single call.** As mentioned earlier, while it is faster to use these batched `attachTagToMultipleObjects()` and `attachMultipleTagsToObject()` calls, there is a limit to the size of a request (about 200KB; refer to section 0) that the tagging service can process. When this limit is exceeded, the request cannot be processed. For `attachTagToMultipleObjects()`, to avoid hitting this limit, we recommend using no more than 2,000 VMs within a single call. For `attachMultipleTagsToObject()`, we recommend no more than 20 tags per call. If you want to associate a single tag with 20,000 VMs, you should do it using 10 iterations with 2,000 VMs per iteration.

In table 3, we summarize the performance of these API calls when attaching 15 tags to 5,000 VMs (75,000 tag associations):

API Call	Number of Iterations	End-to-End Latency
<code>attachTagToMultipleObjects()</code>	45 (15 tags * 3 batches)	6.6 minutes
<code>attachMultipleTagsToObject()</code>	5,000 (one per VM)	1.3 minutes

Table 3: Latency for assigning 15 tags to 5,000 VMs (a total of 75,000 associations). The end-to-end latency depends on the per-association cost for `attachTagToMultipleObjects()` and the per-iteration cost for `attachMultipleTagsToObject()`.



When attaching multiple tags to multiple objects, we recommend using `attachMultipleTagsToObject()`. In addition, we observed that 7.0 U2 outperforms both 6.7 and 7.0 U1 (that is, the end-to-end latencies in 7.0 U2 are lower than in 6.7 and 7.0 U1). As the table indicates, `attachMultipleTagsToObject()` is faster with a larger number of VMs, and overall, the APIs are nearly twice as fast as 6.7. Table 4 shows a few more data points: 10 tags associated with 1,000 VMs, and 1, 10, 15, or 20 (hereafter referred to as {1,10,15,20}) tags associated with 5,000 VMs. We also performed a test where we attached 20 tags to 10,000 VMs. For large numbers of VMs and more than one tag, `attachMultipleTagsToObject()` scales the best. With one tag and many VMs, `attachTagToMultipleObjects()` scales the best.

Operation	10 tags to 1000 VMs (10k associations)	1 tag to 5,000 VMs (5K associations)	10 tags to 5,000 VMs (50K associations)	15 tags to 5,000 VMs (75K associations)	20 tags to 5,000 VMs (100K associations)	20 tags to 10,000 VMs (200K associations)
<code>attach()</code>	3.7 minutes	1.6 minutes	16.1 minutes	23.9 minutes	32.3 minutes	66.2 minutes
<code>attachTagToMultipleObjects()</code>	53 seconds	26 seconds	4.5 minutes	6.9 minutes	9.3 minutes	18.5 minutes
<code>attachMultipleTagsToObject()</code>	16 seconds	1.2 minutes	1.3 minutes	1.3 minutes	2.3 minutes	4.9 minutes

Table 4: End-to-end latency to attach tags to VMs. We show 10 tags associated with 1,000 VMs, {1,10,15,20} tags associated with 5,000 VMs, and 20 tags associated with 10,000 VMs. `attachMultipleTagsToObject()` scales the best when the number of VMs is large.

The *takeaway* points from the table above are as follows:

1. The time to perform a single association increases as more tag associations are added.
2. Batching APIs (`attachTagToMultipleObjects()` or `attachMultipleTagsToObject()`) are preferable to looping over every tag and every VM and using `attach()` on each tag/VM pair.
3. When performing an operation on more than 2,000 VMs or 2,000 tags, break down tags or VMs into batches of less than 2,000. If you don't, you may exceed the maximum request size (200KB) or the maximum response size (7MB). For requests that exceed the limit, you'll see an error such as this: `Hitting vapi endpoint limit during AttachTagToMultipleObjects()`. For responses, the error is `Error: Response size is greater than allowed 7000000b`.

4. As an additional consideration, to avoid denial-of-service attacks to the tagging service, there is a limit of 1500 tagging-related operations per second, increased from 360 in vSphere 6.7 U3. To detect if you are hitting this limit, look for messages like `Request rejected due to high request rate. Try again later.` in the vapi-endpoint log file (`/var/log/vmware/vapi/endpoint/endpoint.log`). If your creation code is exceeding this limit, we recommend inserting a pause of 2 milliseconds or more between tag-creation calls.

## 4.3. Associate tags with other objects

Associating tags with other objects (for example, hosts or datastores or content library items) has the same tradeoffs as associating tags with VMs, so the preceding discussion applies. As with the discussion on VMs, we have 4 cases to consider:

1. To associate 1 tag with 1 object (for example, host), we recommend using `attach(tagID, HostID)`.
2. To associate 1 tag with multiple hosts, we recommend using `attachTagToMultipleObjects(tagID, HostID [])`.
3. To associate multiple tags with 1 host, we recommend using `attachMultipleTagsToObject(tagID[], HostID)`.
4. To associate multiple tags with multiple hosts, we recommend looping over hosts and calling `attachMultipleTagsToObject(tagID[], HostID)` for each host.

The tagging APIs have been optimized for VM and host associations. Thus, most operations might be slower with datastores or other objects vs. VMs and hosts. For example, attaching a single datastore to a single tag might take up to 40 milliseconds, while attaching a single VM to a single tag is typically around 15 milliseconds. The performance difference depends on inventory size and the number of tags, and for most customer setups, the difference will be much lower than what we have indicated above.

### 4.3.1. Query VMs associated with tags

To find the VMs associated with a tag, there are five APIs:

1. `List<TagAssociation.objectToTags> listAttachedObjectsOnTags(List<String> tagIDs)`: Given a list of tagIDs, this call returns the list of objects that are associated with these tags. The return value is a tuple that shows object id (as a DynamicID) and tag id. A DynamicID is an object that contains an identifier and a type. We give an example in the 11. Appendix.
2. `List<DynamicId> listAttachedObjects(tagID)`: Given a single tagID, this call returns a list of all objects (as Dynamic IDs) associated with this tag ID.
3. `List<TagAssociation.tagToObjects> listAttachedTagsOnObjects(List<DynamicID> objectIDs)`: Given a list of VM\_IDs, this function returns all tags associated with those VMs. The return value is a tuple that shows the VM ID and its associated tag IDs.

4. `List<String> listAttachedTags(objectID)`: Given a VM, this call lists all tags associated with the VM.
5. `List<ListResult> list(iterationSpec)` (**Pagination API, new in 7.0**): This call lists all the tag associations that are present in vCenter, provided a page at a time, with a maximum of “limit” per page (default: 5,000). In addition to returning a page of data, this call also returns a marker, which is used to get the next page of data. Please note that this new pagination-based API is accessed via `https://{api-host}/api/vcenter/tagging/associations`, while the other APIs are accessed via `https://{api-host}/rest/com/vmware/cis/tagging`.

Figure 8 below shows the latency to retrieve associated tags or associated objects for 1, 5, 10, or 15 ({1,5,10,15}) tags associated with 20,000 VMs for 4 APIs: `listAttachedObjects()`, `listAttachedObjectsOnTags()`, `listAttachedTagsOnObjects()`, and list using pagination, which is `list()`. In figure 9, we do the same experiment using the `listAttachedTags()` API. We list it separately because its performance is much worse than the other approaches.




Figure 8: End-to-end time to retrieve associations for {1,5,10,15} tags attached to 20,000 VMs using various List APIs.

From figure 8 above, we see the following:

- The `listAttachedObjects()` call takes a single tag as an input. When we have 15 tags, we loop over the 15 tags. Each call returns 20,000 VMs. We have at most 15 iterations (one for each tag). The latency goes from around 2 seconds to 19 seconds, on par with 7.0 U1 and significantly faster than 6.7.

- In contrast, the `listAttachedObjectsOnTags()` call takes a list of tags and returns the objects attached to them. Each call again returns 20,000 VMs. In our experiments, we used 1 tag per call, rather than providing a list of tags, because when we provided 3 or more tags, the call failed due to exceeding request size limits (200KB). We have at most 15 iterations. Because we request tags using the same technique as `listAttachedObjects()`, the latencies are very similar.
- The `listAttachedTagsOnObjects()` call allows us to specify a list of objects. Because of this, we do not have to retrieve the tags for all 20,000 VMs at once. Instead, we break down the list of 20,000 VMs into batches of 2,000 VMs. For 1 tag, we provide 1 tag and iterate over the VMs in batches of 2,000, for a total of 10 iterations. For 15 tags, we provide an array of 15 tags, and again iterate over VMs in groups of 2,000, again for a total of 10 iterations to get associations for 20,000 VMs. This method scales the best, finishing nearly 2x faster than the other approaches. The results are also faster than in 6.7.
- The `list()` (pagination API) call returns associations in a paged manner (for example, if we have 50,000 tag associations in vCenter, the pagination API will get all the tag associations in batches of 5,000 per iteration). In this case, there will be 10 iterations to get all 50,000 tag associations. *Using this API, a client does not need to manually break down the objects into chunks.*




Figure 9: End-to-end time to retrieve associations for {1,5,10,15} tags attached to 20,000 VMs using `listAttachedTags()` API. `listAttachedTags()` is much slower than the other list calls, and we do not recommend using it in environments with more than 2,000 VMs.

From figure 9 above, we see that `listAttachedTags()` is much slower than the other approaches. This call takes only 1 object as an input and returns the appropriate number of attached tags. For 20,000 VMs, we need 20,000 iterations, which is substantially higher than for the other approaches. Although this call is slower than the other approaches, it is significantly faster (3x) in 7.0 U2 vs. 7.0 U1, so for small numbers of VMs, it may be an appropriate choice depending on your requirements.

We summarize the results in table 5 below.

Operation	1 tag associated with 20K VMs	5 tags associated with 20K VMs	10 tags associated with 20K VMs	15 tags associated with 20K VMs
<code>listAttachedObjects</code>	1.5 seconds	7.8 seconds	14.3 seconds	22.8 seconds
<code>listAttachedObjectsOnTags</code>	1.2 seconds	7.2 seconds	13.5 seconds	19.0 seconds
<code>listAttachedTagsOnObjects</code>	3.1 seconds	6.6 seconds	11.5 seconds	15.9 seconds
<code>listAttachedTags</code>	3.3 minutes	3.6 minutes	3.5 minutes	3.5 minutes
<code>List</code> (Pagination)	2.1 seconds	9.1 seconds	18.0 seconds	25.7 seconds

Table 5: Latency for various list APIs. It is best to use APIs like `listAttachedTagsOnObjects()` or the pagination API rather than `listAttachedTags()`.

Because `listAttachedObjectsOnTags()` and `listAttachedObjects()` can return every object in the system, their performance is very sensitive to inventory size. Instead, if it is an option, we recommend using `listAttachedTagsOnObjects()`, since it allows you to break down the number of objects into smaller chunks and give more predictable performance per call. One caveat is that when using `listAttachedTagsOnObjects()`, if any of the input objects has been deleted, the call will return errors for that object, so a client must handle those errors. Alternatively, you can use pagination-based `list()` API, which also gives predictable performance so that you do not have to manually break the object list into chunks of 2,000 entities. In addition, the pagination API will return only valid associations, rather than returning an error if an object is no longer present in the inventory.

It is important to point out that the list APIs above return IDs. Most client applications require names. To retrieve tag names from these IDs, you must use `get()` calls. For example, `Tag.get()` returns a `TagModel` object that has name, id, category\_id, description, and used\_by fields. For performance reasons, if it is practical, we recommend keeping a map of IDs to names on the client-side. Tag names are scoped to categories, so when storing a tag name, be sure to store the category name as well.

As a final example of batching, consider getting the mapping of all VMs associated with each tag. One way to do this is using `listAttachedObjects()`:

```
Get all categories
Get tags for each category
Get VMs associated with Tag using listAttachedObjects()
```

Because the number of objects associated with a tag may exceed 2,000, we recommend using `listAttachedTagsOnObjects()` instead. Because we do not know how many tags will be associated with each object, to avoid overflowing internal vapi-endpoint buffers, one option is to loop over each object individually:

```
Get all objects using vSphere API for retrieving objects (List <DynamicId> objectList)
for (DynamicId object : objectList) {
    List<Tag.objectToTags> result = listAttachedTagsOnObjects(object)
}
```

We could have used `listAttachedTagsOnObjects()`, `listAttachedTags()`, or `list()`, but we chose `listAttachedTagsOnObjects()` (iterating over each object individually) because the return value of `listAttachedTagsOnObjects()` is a mapping of tag to object. If we had used `listAttachedTags()` or `list()`, we would have had to manually create such a mapping.

Here is an example in which we retrieve all associations using the pagination API.

This example also appears in 11.2.13. Example J12: Retrieving all associations using the pagination API.

```
// Get all tag associations a page at a time.
// Construct the tag association IDs.
// Store these IDs in a set.
AssociationsTypes.ListResult page;
AssociationsTypes.IterationSpec iter = new AssociationsTypes.IterationSpec();
static Set<String> associationIds = new HashSet<String>();

iter.setMarker("");
do {
    page = associationsProvider.list(iter);
    iter.setMarker(page.getMarker());

    // Store association IDs for this page
    for (AssociationsTypes.Summary summary : page.getAssociations()) {
        String tagId = summary.getTag();
        String objectId = summary.getObject().getId();
        String associationId = tagId + objectId;
        associationIds.add(associationId);
    }
} while (page.getStatus() != AssociationsTypes.LastIterationStatus.END_OF_DATA);

// Print out the total number of associations
System.out.println("Total number of Tag Associations: " + associationIds.size());
```

The *key takeaway* from the discussion of the list APIs is that two APIs offer predictable performance and are the recommended ways to retrieve associations:

1. `listAttachedTagsOnObjects()`
2. Pagination `list()`

The pagination `list()` API offers robust, predictable performance, since it returns a fixed number of associations per iteration without requiring the client to shard the data into chunks. The pagination `list()` API also retrieves associations only for objects that are currently present in the inventory, so it does not return errors in case an object has been deleted *after* `list()` has been called but *before* associations have been retrieved. The `listAttachedTagsOnObjects()` also offers predictable performance but requires you to properly chunk the list of objects into groups of less than 2,000 objects to avoid hitting limits for the maximum request size (200KB) or maximum response size (7MB). In addition, if the list of objects includes deleted objects, then `listAttachedTagsOnObjects()` will return errors for those objects, so a client must handle those errors.

## 5. Tagging APIs: PowerShell

The previous examples all used Java. In this section, we discuss using PowerShell. As described in [Writing Performant Tagging Code: Tips and Tricks for PowerCLI](#) [1], there are typically two ways to write PowerShell scripts for vSphere tags: either using cmdlets or using direct tagging APIs. For performance reasons, we recommend using direct tagging APIs instead of cmdlets, as indicated in the blog article.

The direct tagging service calls in PowerShell use the same vCenter server and tagging service resources as Java. Thus, the performance results are the same as the Java results above. In this section, we focus on the difference in the API calls themselves. The API calls differ slightly from Java to PowerShell. Java uses camelCase, while PowerShell puts underscores in the method name. We give some examples below in table 6.

Description	Java	PowerShell (CIS service)
Attach one tag to one VM	<code>Attach()</code>	<code>Attach()</code>
Attach one tag to multiple objects	<code>attachTagToMultipleObjects()</code>	<code>Attach_tag_to_multiple_objects()</code>
Attach multiple tags to one object	<code>attachMultipleTagsToObject()</code>	<code>Attach_multiple_tags_to_object()</code>
List objects attached to a tag	<code>listAttachedObjectsOnTags()</code>	<code>List_attached_objects_on_tags()</code>
List associations (pagination)	<code>List()</code>	<code>List()</code> (instead of <code>com.vmware.cis</code> , please use <code>com.vmware.vcenter</code> )

Table 6: Examples of differences between Java and PowerShell method names. Note that the pagination API for PowerShell is at a different location from the other association-based APIs, using `com.vmware.vcenter` instead of `com.vmware.cis`.

There are also some minor differences in how arrays and objects are created between PowerShell and Java. Please see the examples in 11. Appendix for details.

PowerCLI version 11.5.0 is compatible with vCenter 6.0 through 7.0 U1. At the time of this writing, PowerCLI version 12.0 up to 12.7.0 is compatible with vCenter 7.0 U2. Also refer to the [VMware PowerCLI documentation](#) [6].



## 6. Linked Mode

### 6.1. Replication Performance

In [Enhanced Linked Mode \(ELM\)](#) [7], vCenter servers replicate tag and category information between peer nodes periodically. Each node replicates data with its peer every 30 seconds, so changes to categories and tag definitions on one node may take some time before they propagate to the other nodes in the ELM setup. The choice of topology can impact this replication time. For example, consider figure 10, which shows a [simple ring topology](#) [9] for vCenter servers.




Figure 10: Ring topology for linked mode

In the figure above, the replication partners are as follows:

- VC 1: VC 2 and VC 6
- VC 2: VC 1 and VC 3
- VC 3: VC 2 and VC 4
- VC 4: VC 3 and VC 5
- VC 5: VC 4 and VC 6
- VC 6: VC 1 and VC 5

When a new tag is added on VC 1, here is a rough timeline of replication:

1. Time 0: tag added to VC 1
2. Up to 30 seconds later: VC 1 replicates to VC 2 and VC 6
3. Up to 30 seconds later: VC 2 replicates to VC 3; VC 6 replicates to VC 5
4. Up to 30 seconds later: VC 3 replicates to VC 4; VC 5 also replicates to VC 4, but ELM performs conflict resolution to ensure only one new tag definition is created on VC 4

As the timeline above suggests, in this 6-node ring topology, a tag may take up to 90 seconds to replicate throughout the system. Because each node pushes changes at 30-second intervals, this is the worst-case latency. The nodes do not push changes at the same time, so in the best case, a tag

change may occur right before that node has been scheduled to push changes. In this case, the changes would be reflected more quickly.

Different topologies will have different delays. For example, suppose VC 1 was NOT connected to VC 6, as shown below in figure 11.




Figure 11: Linear topology for linked mode. We do not recommend this approach due to latency of tag/category replication.

1. Time 0: tag added to VC 1
2. Up to 30 seconds later: VC 1 replicates to VC 2
3. Up to 30 seconds later: VC 2 replicates to VC 3
4. Up to 30 seconds later: VC 3 replicates to VC 4
5. Up to 30 seconds later: VC 4 replicates to VC 5
6. Up to 30 seconds later: VC 5 replicates to VC 6.

As the timeline suggests, the worst-case latency for the last node to receive a tag update is 30 seconds \* 5 = 150 seconds after the tag has been added. The best case is nearly instantaneous if the changes occur right before a node is scheduled to push tag updates.

## 6.2. Comparison to standalone vCenter tagging performance

In Enhanced Linked mode, tag and category definitions (for example, Category “VM OS” or Tag “Linux”) are global to all vCenters, while tag associations (for example, VM-1 is a “Linux” VM) are local to each vCenter, since a given VM, host, datastore, or other entity can reside only on one vCenter. Because the associations do not need to be replicated, most tagging operations can proceed completely in parallel. To illustrate this point, we constructed a 5-node ring topology like the one described in figure 10. Each vCenter had 5,000 VMs, so for the 5 vCenters, we had a total of 25,000 VMs. Using the `AttachMultipleTagsToObject()` API, we associated 20 tags with each VM. This means that on each vCenter, we would be creating 20 tags \* 5,000 VMs = 100,000 associations for a total of 500,000 associations overall across the entire linked mode installation. In section 0, we show that `AttachMultipleTagsToObject()` with 75,000 associations takes around 1.3 minutes for an individual vCenter, so we would expect 100,000 associations to be slightly longer. As shown below in table 6 for a standalone vCenter, the time is approximately 90 seconds, in line with expectations. When we run the same workload (attach 20 tags to each VM) against each of

the vCenters in a 5-vCenter Enhanced Linked mode topology, with each vCenter managing 5,000 VMs, we expect each vCenter to complete in the same amount of time as the standalone time, since there are no dependencies between vCenters and they can execute completely in parallel. As table 6 indicates, the overall end-to-end time for each vCenter is indeed close to 90 seconds. In this case, with 5 vCenters, each with 100,000 associations, a 5-vCenter Enhanced Linked Mode topology can handle up to 500,000 associations.

Configuration	Latency to attach 20 tags to 5,000 VMs (in seconds)
Standalone vCenter	90
Enhanced Linked Mode: vCenter 1	92
Enhanced Linked Mode: vCenter 2	88
Enhanced Linked Mode: vCenter 3	91
Enhanced Linked Mode: vCenter 4	98
Enhanced Linked Mode: vCenter 5	90

Table 6: Latency to attach 20 tags to 5,000 VMs, standalone vCenter vs. vCenters in an Enhanced Linked Mode configuration

## 7. Tags vs. Custom Attributes

There are some situations in which tags might not be the right fit. For example, suppose a user has 4,000 VMs in their vSphere environment, and they wish to assign an assetID to each VM. One option is to create a category called “assetID,” and then have each tag within that category be a different assetID. However, this approach would generate numerous tags, which may cause performance issues. Another example might be assigning a “last power-off time” to each VM. One could have a category named “last power-off time” and have each tag be a different time, but this can again cause scalability issues. For situations like these two cases, it might be better to use vSphere custom fields and custom attributes. A discussion of custom fields and custom attributes is beyond the scope of this paper: please refer to official [VMware documentation](#) [9] or [William Lam’s blog on this topic](#) [11].

## 8. Scale Numbers for Good Performance

As mentioned in 2. Introduction, there are no hardcoded limits for the number of categories, tags, or tag associations that the tagging service can support. However, as we have shown in this paper, performance of the UI and APIs depends on the number of associations.

Though primarily an upgrade release, 7.0 U2 adds some performance improvements to the tagging infrastructure. In table 7, we list scale numbers that will work well *out-of-the-box* for vSphere 7.0 U2. The limits are the same for 7.0 U1 and increased relative to 6.7: the supported number of tag associations is 150,000 per vCenter for 7.0 U1 and 7.0 U2 vs. 25,000 for 6.7. For these numbers of tags, categories, and associations, and assuming proper hardware configuration of vCenter (see 11. Appendix), we expect acceptable out-of-the-box performance for the UI, second- and third-party applications, and custom scripts written by customers, as long as the scripts adhere to our best practices, the system has appropriate hardware, and the VCSA is properly sized (in our tests, we used the X-Large configuration of the VCSA). As mentioned earlier, tag definitions and category definitions are shared among all nodes in a Linked-mode domain, while associations are scoped to a single vCenter. Thus, the associations below are for a single vCenter Server in a Linked-mode domain, and the number of tag and category definitions are the cumulative number of definitions across all nodes in the Linked-mode domain. The number of associations scales with the number of nodes in Linked mode, so if you have, say, 10 nodes in Linked mode, you can support up to 1,500,000 total associations across those nodes, with no more than 150,000 per vCenter (in such a situation, each VCSA would need to be an X-Large configuration).

Target of Limit	Limit
Category definitions	6,000
Tag definitions	8,000
Tag Associations	150,000

Table 7: vSphere 7.0 U2 tag limits for acceptable out-of-the-box performance. The number of associations is for a single vCenter in a Linked Mode configuration.

The limits above are based on response latency as well as hardcoded limits within vSphere processes to prevent denial-of-service (DoS) attacks and out-of-memory exceptions for large response bodies. As mentioned above, the limits are suggested (soft) limits: with careful coding and proper VCSA sizing, as indicated in this paper, you can exceed these limits.

## 9. Concluding Remarks

In this paper, we have provided information about the proper use of tagging APIs for performance. We have discussed the following considerations:

### Coding tips:

- For faster tag or category creation time, consider using multiple threads.
- If possible, for better performance, create tags spread across many categories, rather than putting all tags in a few categories.
- For faster assignment of tags to objects, consider using `attachMultipleTagsToObject()` or `attachTagToMultipleObjects()` instead of `attach()`.
- For faster listing of tag associations, consider `listAttachedObjectsOnTags()`, `listAttachedTagsOnObjects()`, or pagination `list()` API, rather than `listAttachedObjects()` or `listAttachedTags()`. The first two calls allow the user to specify a concrete list of tags or objects, and they allow chunking the requests to avoid overwhelming the tagging service. The `list()` API is especially helpful if you prefer not to manually chunk the objects on the client side. It is also easier from an error-handling perspective, as described earlier in the text.
- Use a ring-based ELM topology for best latency in replicating tags. For more details, refer to [vCenter Server Architecture Part 1 – The Basics](#) [8]
- Keep a local map of IDs (tag IDs, category IDs, VM\_IDs) to names so that you don't have to make a call to vCenter or the tagging service every time you need an ID from a name.

### Important caveats:

- For attach calls, the maximum request size we recommend is 2,000 VMs or 2,000 tags to stay below the maximum request size supported by the vapi-endpoint process. If possible, consider these limits when choosing which APIs calls to use.
- For listing associations, if using `listAttachedTagsOnObjects()`, be sure to break down the list of objects into groups of 2,000.
- Consider using shorter tag names if you need to exceed 2,000 tags per attach call. The limit of 2,000 is based on the number of bytes in the response, not the sheer number of tags, so smaller tag names may allow you to retrieve more items in a single call.

- The maximum response size that can be returned by tagging APIs is approximately 7MB. In practice, for 256 billion tags, this limit represents a response size of approximately 30,000 tag associations: this is, for example, 2,000 objects and 15 tags per object, or 200 objects and 150 tags per object, or 1,000 tags and 30 tags per object. When using `listAttachedTagsOnObjects()` or `listAttachedObjectsOnTags()` to return tag associations, be careful to avoid reaching this limit by chunking the request sizes by tags or by objects. The error you receive if you exceed this limit is shown in 4.2.2.1. Chunking VMs in calls to `attachTagToMultipleObjects()`. Note that the pagination `list()` API automatically breaks down responses to avoid this limit. Creating and deleting tags creates tombstones in the tagging service database. These tombstones are needed for proper replication in linked mode, but excessive numbers of them can impact performance. Please see [KB 52387 \[5\]](#).
- There are internal vCenter limits to avoid DoS attacks. The vCenter appliance allows no more than 1,500 requests per second to tagging and other RESTful vAPI-based services within the VCSA (for example, content library). If you are exceeding this limit in your environment, we recommend inserting a pause periodically to slow down the rate of requests. To know if you are hitting this limit, please check the vAPI-endpoint log file in `/var/log/vmware/vapi/endpoint/endpoint.log`.
- In certain rare cases, you may need to increase the heap size of the vAPI-endpoint service to accommodate high amounts of tagging traffic. When the vAPI-endpoint approaches its heap size, it triggers an alert within vCenter, so please be careful to check for this alert periodically.
- The latency to perform a tag association depends on the number of pre-existing associations. If many associations exist, it may be necessary to perform maintenance on the tag association table, which is a table within the main vCenter database. To perform such cleanup, please run the following vacuum job to clean up the database engine statistics:
  - a. Connect to your VCSA VM via SSH as the root user.
  - b. Log into postgres:
 

```
/opt/vmware/vpostgres/current/bin/psql -U postgres -d VCDB
```
  - c. Run vacuum analyze in this order:
 

```
vacuum analyze cis_kv_keyvalue;
vacuum analyze cis_kv_providers;
```

## 10. References

- [1] VMware. (2020) VMware vSphere Automation SDKs Programming Guide. <https://vdc-download.vmware.com/vmwb-repository/dcr-public/2fa085a7-3ff5-477e-9cf9-5a03eff7549a/68a01bae-01c6-4d48-84b5-e7678d1921dc/vsphere-automation-sdk-701-programming-guide.pdf>
  
- [2] VMware. (2020) vSphere Automation SDK for Java. <https://techpartnerhub.vmware.com/sdks/vsphere-automation-java/7.0>
  
- [3] Ravi Soundararajan and Joseph Zuk. (2019, June) Writing Performant Tagging Code: Tips and Tricks for PowerCLI. <https://blogs.vmware.com/performance/2019/06/writing-performant-tagging-code-tips-and-tricks-for-powercli.html>
  
- [4] Raju Angani, Ravi Soundararajan, and Maarten Wiggers. (2020, August) VMware vSphere 6.7 Tagging Best Practices. <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/performance/tagging-vsphere67-perf.pdf>
  
- [5] VMware. (2018, November) Troubleshooting and addressing accumulation of tombstones in a Platform Services Controller (52387). <https://kb.vmware.com/s/article/52387>
  
- [6] VMware. PowerCLI SDK. <https://code.vmware.com/tool/vmware-powercli>
  
- [7] VMware. (2019, May) vCenter Enhanced Linked Mode for a vCenter Server Appliance with Embedded Platform Services Controller. <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vcenter.install.doc/GUID-7A90361C-E95B-4C53-A328-3DB9EE20536D.html>
  
- [8] Emad Younis. (2018, May) vCenter Server Architecture Part 1 – The Basics. <https://emadyounis.com/vcenter-server-architecture-part-1-the-basics/>

- [9] VMware. (2019, May) Custom Attributes in the vSphere Web Client. <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vcenterhost.doc/GUID-73606C4C-763C-4E27-A1DA-032E4C46219D.html>
- [10] William Lam. (2015, January) Custom Attributes != vSphere Tags. <https://www.virtuallyghetto.com/2015/01/custom-attributes-vsphere-tags.html>
- [11] VMware. VMware vSphere Automation SDK for Java. <https://github.com/vmware/vsphere-automation-sdk-java>
- [12] VMware. Package com.vmware.cis.tagging. <https://vmware.github.io/vsphere-automation-sdk-java/vsphere/6.7.1/vsphereautomation-client-sdk/com/vmware/cis/tagging/package-summary.html>
- [13] VMware. Tagging Service SDK (PowerCLI). <https://code.vmware.com/docs/7335/powercli-11-0-0-user-s-guide/doc/GUID-F0088D49-CCFC-4F26-9C8C-845DE5AA951A.html?h=tags>
- [14] VMware. vSphere Java APIs. <https://github.com/vmware/vsphere-automation-sdk-java>
- [15] VMware. Java API for vSphere 7.0. <https://code.vmware.com/web/sdk/7.0/vsphere-automation-java>



# 11. Appendix

## 11.1. Experimental setup and software versions

We used VMware PowerCLI 10.1.0 build 8346946. Note that newer versions may demonstrate better performance.

We used Java 1.8 (“1.8.0\_231”).

Our vCenter was vCenter 7.0 U2 running on ESXi 6.5 U2. Our inventory was composed of 2,000 hosts and 35,000 VMs (simulated using a proprietary VMware-internal tool). Because of the large number of VMs, we used the “Extra Large” size for the VCSA.

## 11.2. Java code examples

Some of these examples come from the following Github repository:

<https://github.com/vmware/vsphere-automation-sdk-java> [12].

For each of the Java examples below, we need a set of imports. Here are the imports we used in our code snippets:

```
import vim25
import com.vmware.cis.tagging.Category;
import com.vmware.cis.tagging.CategoryTypes;
import com.vmware.cis.tagging.Tag;
import com.vmware.cis.tagging.TagAssociation;
import com.vmware.cis.tagging.TagModel;
import com.vmware.cis.tagging.TagTypes;
import com.vmware.cis.tagging.CategoryModel.Cardinality;
```

### 11.2.1. Example J0: VM DynamicID and Tag TagID

Many code samples involve the use of VM DynamicID objects. A DynamicID object is similar to a Managed Object Reference (MoRef, described in [10]), but the identifier and type are stored in separate fields. Here is a simple code snippet for constructing a DynamicID from a VM MoRef. We first retrieve the VM by name, and then we get its MoRef. Finally, we construct a DynamicID from this MoRef.

```
// Retrieve the VM MoRef from its name
this.VMMoRef = VimUtil.getVM(this.vimAuthHelper.getVimPort(),
                           this.vimAuthHelper.getServiceContent(),
                           "VM_NAME");
assert this.VMMoRef != null;
// convert the MoRef to vAPI DynamicID
this.vmDynamicId = new DynamicID(this.VMMoRef.getType(), this.VMMoRef.getValue());
```

### 11.2.2. Example J1: Retrieving category/tag IDs from category/tag names

Most tagging APIs require tag IDs. A tag ID is globally unique, while a tag name is unique within a category. To find a tag ID given a tag name, you must also specify the category ID.

The following code finds a category ID given a category name:

```
Category _categoryService = vapiAuthHelper.getStubFactory().createStub(Category.class,
                                                                    sessionStubConfig);

// Get CategoryId
public String getCategoryId(String catName) {
    List<String> catIds = _categoryService.list();
    String retCatId = NULL;
    if (catIds.size() > 0) {
        for (String cat : catIds){
            if (_categoryService.get(cat).getName().equals(catName)) {
                retCatId = _categoryService.get(cat).getId();
                break;
            }
        }
    }
    return retCatId;
}
```

The following code gets a tag ID given a tag name and a category name, using the routine above to map the category name to an id.

```
Tag _taggingClient = vapiAuthHelper.getStubFactory().createStub(Tag.class,
                                                                sessionStubConfig);

// Get TagId for a given tagName and CategoryName
public String getTagId(String tagName, String catName) {
    List<String> tagIds = _tagProvider.list();
    String catId = getCategoryId(catName);
    String retTagId = "";
    if (tagIds.size() > 0) {
        for (String tag : tagIds) {
            if (_tagProvider.get(tag).getName().equals(tagName)
                && _tagProvider.get(tag).getCategoryId().equals(catId)) {
                retTagId = _tagProvider.get(tag).getId();
                break;
            }
        }
    }
    return retTagId;
}
```

### 11.2.3. Example J2: Create Category

```
/**
 * Creates a tag category
 *
 */
private String createTagCategory(String name, String description,
                                Cardinality cardinality, Set<String> AssociableTypes) {
    CategoryTypes.CreateSpec createSpec = new CategoryTypes.CreateSpec();
    createSpec.setName(name);
    createSpec.setDescription(description);
    createSpec.setCardinality(cardinality);
    createSpec.setAssociableTypes(associableTypes);
    return this.categoryService.create(createSpec);
}
// Sample invocation: create category named "GuestOS"
Set<String> associableTypes = new HashSet<String>(); // empty hash set
associableTypes.add("VirtualMachine");
String categoryId = createTagCategory("GuestOS",
                                     "Guest OS",
                                     CategoryModel.Cardinality.SINGLE,
                                     associableTypes);
```

### 11.2.4. Example J3: Create Tag

In this example, we use the categoryId from 11.2.3. Example J2: Create Category above.

```
/**
 * Creates a tag
 *
 */
private String createTag(String name, String description,
                        String categoryId) {
    TagTypes.CreateSpec spec = new TagTypes.CreateSpec();
    spec.setName(name);
    spec.setDescription(description);
    spec.setCategoryId(categoryId);

    return this.taggingService.create(spec);
}
// Sample invocation: create tag named "Windows 10" in categoryId from above
String tagID = createTag("Windows 10", "DB Operating System", categoryId);
```

### 11.2.5. Example J4: Associating a tag with a VM (attach())

In this example, we use the VM DynamicID (vmDynamicId) from “11.2.1. Example J0: VM DynamicID and Tag TagID“ above and the tagID from “11.2.4. Example J3: Create Tag” above.

```
/**
 * Associates tagId from Example J3 to VM whose DynamicID is vmDynamicId (Example J0)
 * We assume vmDynamicId is retrieved as shown in Example J0 above.
 *
 */

// Get a handle to the tagAssociation Methods
tagAssociation = vapiAuthHelper.getStubFactory().createStub(TagAssociation.class,
                                                         sessionStubConfig);

tagAssociation.attach(tagId,vmDynamicId);
```

### 11.2.6. Example J5: Associating a tag with multiple VMs (attachTagToMultipleObjects())

Attach 1 tag to 1,000 objects

```
// Attach 1 tag to first 1,000 VMs
// Assume we have a list of VM DynamicIDs → List<DynamicID> vmDynamicIds

int totalVMs = 1000;
List<DynamicID> vmDynamicIdsNeeded = new LinkedList<>();
for (int k = 0; k < totalVMs; k++) {
    vmDynamicIdsNeeded.add(vmDynamicIds.get(k));
}

// Get a handle to the tagAssociation Methods
tagAssociation = vapiAuthHelper.getStubFactory().createStub(TagAssociation.class,
                                                         sessionStubConfig);
tagAssociation.attachTagToMultipleObjects(tagId,vmDynamicIdsNeeded);
```

### 11.2.7. Example J6: Associating multiple tags with multiple VMs (attachMultipleTagsToObject())

Attach 10 tags to 1,000 objects

```
// Attach 10 tags to 1,000 VMs
// Assumes we have a list of all tagIDs → List<String> tagIds
// Assumes we have a list of all VM_IDs → List<DynamicID> vmDynamicIds
// We will assign the first 10 of these tags to 1,000 VMs

int totalVMs = 1000;
int totalTags = 10;

// Get first 10 tags from the total tagIds
List<String> associableTags = new LinkedList<>();
for (int k = 0; k < totalTags; k++) {
    associableTags.add(tagIds.get(k));
}

// Get first 1,000 VMs from the total VM_IDs
List<DynamicID> vmDynamicIdsNeeded = new LinkedList<>();
for (int k = 0; k < totalVMs; k++) {
    vmDynamicIdsNeeded.add(vmDynamicIds.get(k));
}

tagAssociation = vapiAuthHelper.getStubFactory().createStub(TagAssociation.class,
                                                         sessionStubConfig);
for (int i = 0; i < totalVMs; i++) {
    tagAssociation.attachMultipleTagsToObject(vmDynamicIdsNeeded.get(i), TagIds);
}
```

### 11.2.8. Example J7: List tags associated with a VM

```
// List tags associated with vm vmDynamicId from above.
// Assumes tagAssociation object from above
List<String> retTagIds = tagAssociation.listAttachedTags(vmDynamicId);
```

### 11.2.9. Example J8: List VMs associated with a tag

```
// List tags associated with tagId from above
// Assumes tagAssociation object from above
List<DynamicID> retDynamicIds = tagAssociation.listAttachedObjects(tagId);
```

### 11.2.10. Example J9: List tags associated with a group of VMs

```
// Helper object for storing results
ObjectToTags {
    DynamicID objectId;
    List<String> tagId;
}

// List tags associated with 10 VMs from vmDynamicIdsUsed list above
// Assumes tagAssociation object from above
List<TagAssociationTypes.ObjectToTags> retObjectToTags =
    tagAssociation.listAttachedTagsOnObjects(vmDynamicIdsUsed.subList(0,10));
```

### 11.2.11. Example J10: List VMs associated with a group of tags

```
// Helper object for storing results
TagToObjects {
    String tagId;
    List<DynamicID> objectIds;
}
// Assumes we have a list of all tagIDs → List<String> tagIds
List<TagAssociationTypes.TagToObjects> retTagToObjects =
    tagAssociation.listAttachedObjectsOnTags(tagIds);
```

### 11.2.12. Example J11: Retrieving tag information (e.g., name, description) from the tag ID

```
// TagModel contains name, description, tag ID, category ID, and usedBy fields
TagModel retTagModel = _tagProvider.get(tagID);
```

### 11.2.13. Example J12: Retrieving all associations using the pagination API

Here is an example in which we retrieve all associations using the pagination API.

```
// Get all tag associations a page at a time.
// Construct the tag association IDs.
// Store these IDs in a set.
AssociationsTypes.ListResult page;
AssociationsTypes.IterationSpec iter = new AssociationsTypes.IterationSpec();
static Set<String> associationIds = new HashSet<String>();

iter.setMarker("");
do {
    page = associationsProvider.list(iter);
    iter.setMarker(page.getMarker());

    // Store association IDs for this page
    for (AssociationsTypes.Summary summary : page.getAssociations()) {
        String tagId = summary.getTag();
        String objectId = summary.getObject().getId();
        String associationId = tagId + objectId;
        associationIds.add(associationId);
    }
} while (page.getStatus() != AssociationsTypes.LastIterationStatus.END_OF_DATA);

// Print out the total number of associations
System.out.println("Total number of Tag Associations: " + associationIds.size());
```

## 11.3. PowerShell code examples

In this section, we give select examples of PowerShell. Examples of creating a category, creating a tag, `attach()`, `attach_tag_to_multiple_objects()`, and `attach_multiple_tags_to_object()` are given in the [Performance Team blog on tagging](#) [1] and are therefore not repeated here.

### 11.3.1. Example P0: Creating VMID Objects and retrieving Tag IDs

```
# Retrieving a tag ID from its name
function Get-TagIdFromName {
    Param ($a)
    $allTagMethodSvc = Get-CisService com.vmware.cis.tagging.tag
    $allTags = @()
    $allTags = $allTagMethodSvc.list()
    foreach ($tag in $allTags) {
        if (($allTagMethodSvc.get($tag.value)).name -eq $a) {
            return $tag.value
        }
    }
    return $null
}

# Creating a single VM ID Object
$testVM = Get-VM -Name "myVM"

$VMInfo = $testVM.ExtensionData.MoRef
$vmid = New-Object PSObject -Property @{
    id = $VMInfo.value
    type = $VMInfo.Type
}

# Create an array of VM IDs
$allVMs = Get-VM
$useTheseVMIDs = @()

# Create VM objects for all VMs.
# This should be done once for all VMs, not every time
# you want to do an association.
for ($i = 0; $i -lt 1000; $i++) {
    $VMInfo = $allVMs[$i].extensiondata.moref
    $useTheseVMIDs += New-Object PSObject -Property @{
        id = $VMInfo.value
        type = $VMInfo.type
    }
}
}
```



### 11.3.2. Example P1: List tags associated with a VM

```
$allCategoryMethodSVC = Get-CisService com.vmware.cis.tagging.category
$alltagMethodSVC = Get-CisService com.vmware.cis.tagging.tag
$allTagAssociationMethodSVC = Get-CisService com.vmware.cis.tagging.tag_association

# Use first VM from above list
$useThisVMID = $useTheseVMIDs[0]
$tagList = $allTagAssociationMethodSVC.list_attached_tags($useThisVMID)
```

### 11.3.3. Example P2: List VMs associated with a tag

```
$allTagMethodSvc = Get-CisService com.vmware.cis.tagging.tag

# Use method above to get tag ID from tag name
$tagId = Get-TagIdFromName "myTag"
$vmList = $allTagAssociationMethodSVC.list_attached_objects($tagId)
```

### 11.3.4. Example P3: List tags associated with a group of VMs

```
$assocSVC = Get-CisService com.vmware.cis.tagging.tag_association

# Pass in VMs in groups of 2,000
$sampleSize = 2000

# Assume a list of VMs: $useTheseVMIDs
$index = 0
$tagList = @()
for ($i = 0; $i -le ($useTheseVMIDs.Count/$sampleSize); $i++) {
    $tagList += $assocSVC.list_attached_tags_on_objects($useTheseVMIDs[$index..($index +
$sampleSize)])
}
```

### 11.3.5. Example P4: List VMs associated with a group of tags

```
$assocSVC = Get-CisService com.vmware.cis.tagging.tag_associations

# Loop over 15 tags and get the result.
# If there are too many VMs and you pass in all 15 tags in one go, you may get a RESPONSE
error
$tagIdList = $tagList[0 .. 15]
$VMAssocList = @()
foreach ($tag in $tagIdList){
    # Must use a list because list_attached_objects_on_tags assumes a list of tags as an input
    # We are doing this one at a time, but you can use multiple at a time as long
    # as you do not get a RESPONSE error
    $useTheseTags = @()
    $useTheseTags += $tag
    $VMAssocList += $assocSVC.list_attached_objects_on_tags($useTheseTags)
}
```

### 11.3.6. Example P5: Retrieve Associations using Pagination API

```
# Get a connection to the new tagging API (not: com.vmware.vcenter, not com.vmware.cis)
$assocSVC = Get-CisService com.vmware.vcenter.tagging.associations

# Create an object to help us iterate through the list of associations
$iterator = $assocSVC.Help.list.iterate.Create()

# Loop over pages until we have received all associations
# When the result status is END_OF_DATA, we have received all pages
do {
    $result = $assocSVC.list($iterator)

    # display associations
    $result.associations

    # set the iteration marker to the next page
    $iterator.marker = $result.marker
} while ($result.status -ne "END_OF_DATA")
```

## About the Authors

**Raju Angani** is a staff performance engineer in the Performance team at VMware. His main areas of focus are vSphere tagging performance and vCenter database performance.

**Ravi Soundararajan** is a principal engineer in the Performance team at VMware, specializing in vCenter performance and scalability. He has presented numerous times about vCenter at VMworld. His twitter handle is @vCenterPerfGuy. Ravi received his SB from MIT and his MS and PhD degrees from Stanford, all in Electrical Engineering.

**Maarten Wiggers** is a senior staff engineer and works on vCenter, with a focus on tagging, policies, and resource management. He has been with VMware for 8 years and has a PhD from the University of Twente, The Netherlands.

## Acknowledgements

We gratefully acknowledge Joseph Zuk, Oleg Zaydman, Abhijith Marulaiah Prabhudev, Chung-Yen Chang, Udaya Ganiga, and Dinesh Suresh for their assistance with the content of this paper.