

vSphere Web Services SDK 6.7 Release Notes

Release date: 17 APR 2018 | ESXi build 8169922, SDK build 7970399

For vSphere 6.7 GA. Last document update: 24 APR 2018

Check frequently for additions and updates to these release notes.

Overview

The vSphere Web Services SDK supports the development of vSphere client applications that use the vSphere API to manage virtual machines and other virtual infrastructure components, such as datacenters, datastores, and networks.

Distribution Kit

The vSphere Web Services SDK is distributed as part of the vSphere Management SDK. The vSphere Management SDK is a collection of vSphere SDKs. When you extract the contents of the distribution kit, the vSphere WSSDK is contained in the SDK sub-directory:

```
VMware-vSphere-SDK-6.7.0-buildNumber.zip
```

```
SDK
  eam
  sms-sdk
  spbm
  ssoclient
  vsphere-ws
```

New in This Release

- **New vSphere APIs.**

For information about additions to the vSphere API (new managed object types, data object types, properties, and methods), see the *vSphere API Reference*. The main landing page for the *vSphere API Reference* contains links to this information in the section “What’s New in vSphere API 6.7?”

- **New `InstantClone` method added.**

The vSphere 6.7 release includes the `InstantClone_Task()` method for low-latency deployment of cloned virtual machines from a running or frozen source virtual machine.

- **New Enhanced vMotion Compatibility feature added.**

The vSphere 6.7 release includes the Enhanced vMotion Compatibility (EVC) feature. EVC facilitates vMotion between different CPU generations through use of Intel Flex Migration and AMD-V Extended Migration technologies. When enabled on a particular virtual machine, the EVC feature ensures that CPU compatibility retains the specified mode regardless of migration across clusters, datacenters, or vCenter server instances. The SDK includes a new sample program, `VMApplyEvc.java`, that demonstrates how to use the new feature. See the documentation for `applyEvcModeVMTask()` in the *vSphere API Reference* for reference information.

- **.NET SDK DLLs added.**

The SDK now includes pre-built DLLs for .NET, so there is no need for you to build the C# DLLs before running the samples.

Deprecations

Several properties and types defined for the API have been deprecated as of vSphere API 6.7. You can see all of the objects, methods, and parameters with indicators to show when they were added and/or deprecated by reviewing the API Versions Reference table in the *vSphere Web Services API Reference*.

Known Issues

The following section describes the known issues for the vSphere Web Services SDK in release 6.7.

- **While acquiring bearer token, Java samples fail with SOAP exception.**

When JAX-WS sample programs are run with Java 1.8.0_161, 1.8.0_171/172, or 10.0.1, they throw the error “Entity References are not allowed in SOAP documents” after trying to acquire a bearer token, and exit with a back-trace. One solution is to use the more stable JDK 1.8.0_131.

- **JDK 1.8 Update 141 and 144 SSO connection failure.**

JDK 1.8 Update 141 and 144 would fail to complete an SSO connection to vCenter Server. The vpxd log would show the following lines:

```
ERROR task-internal... vim.SessionManager.loginByToken: vim.fault.InvalidLogin:
--> Result:
--> (vim.fault.InvalidLogin) {
--> faultCause = (vmodl.MethodFault) null,
--> msg = ""
--> }
```

For information about workarounds, see [Knowledge Base article 2151706](#).

- **.NET samples fail to access vCenter Server by using an IPv6 address.**

If you use an IPv6 address to access the vCenter Server system, the samples fail to connect to the vSphere services and the following error is logged:

```
System.ServiceModel.EndpointNotFoundException: There was no endpoint listening at https://[fd01:0:101:2605:0:a:0:1a61]/sts/STSService that could accept the message. This is often caused by an incorrect address or SOAP action. See InnerException, if present, for more details. --> System.Net.WebException: The remote server returned an error: (503) Server Unavailable.
```

Workaround: Use an FQDN to access the vSphere services when running the .NET samples.

- **FindByDnsName case sensitivity.**

The FindByDnsName operation expects a `dnsName` argument to specify a target domain name. For requests to vCenter Server, the argument value is case insensitive. For requests to ESX hosts, the argument value is case sensitive.

- **HTTP file download performance.**

HTTP file download performance using a vCenter client connection takes longer than the same file transfer using an ESXi connection. A similar issue with HTTP file upload has been resolved.

- **SCSI LUN queue depth is not available.**

The ESXi Server does not set the `ScsiLun.queueDepth` property. Modifications to this property will have no effect. Property retrieval will not produce a value.

- **Disk chains of managed virtual disks not fully supported.**

In this release, vSphere does not fully support managed virtual disk chains.

- When you delete a managed virtual disk chain, only the leaf disk file is deleted. The leaf disk is the last child disk in the chain, while the base disk is the earliest parent in the chain.
- If your application manages disk chains that span datastores, the application must maintain its own mapping between child disks and datastores. You need the mapping to delete the entire disk chain correctly.
- If your application manages disk chains that span datastores, certain situations can cause disk IDs to behave unexpectedly. vSphere maintains an association between disk ID and the file path of the leaf disk. When a child disk is created, the file path for the ID in the managed virtual disk inventory is updated to point to the newly created leaf disk.

If you subsequently use an API method to reconcile a different datastore inventory that contains the base disk of the chain, or one of the parent disks in the interior of the chain, vSphere creates another association of the disk ID that points to the parent disk on that datastore. If you reconcile all datastores, vSphere creates another association with the disk ID for each datastore that backs the chain. For this reason, your application needs to keep track of the datastores, especially the datastore that contains the leaf disk of the chain.

Workaround: To avoid creating duplicate disk ID associations for a chain of managed virtual disks, use a single datastore as a backing resource for all disks in a chain.

If you already have a chain that spans multiple datastores, you can create a copy of the chain on a single datastore by using the `CloneVStorageObject_Task()` method of the VMware Web Services API. The parameters to the clone method must supply the datastore of the leaf disk as the source datastore. The datastore parameter resolves any ambiguity with the disk ID associations.

If you need to verify the properties of a managed virtual disk ID, the associated datastore is uniquely identified by the property `RetrieveVStorageObject().config.backing.datastore.summary.url`, and the associated virtual disk file is located by the property `RetrieveVStorageObject().config.backing.filePath`.

After you clone the disk chain, you can delete the old chain that spanned datastores. Deleting a managed virtual disk chain is described in the KB article [2147767](#).

Documentation Issues

- **VirtualMachine.CloneVM_Task method failure.**

The *vSphere Web Services API Reference* does not document the following limitation with regard to `VirtualMachineCloneSpec`:

When the `VirtualMachine.CloneVM_Task` method fails, it produces the following error, "The specified delta disk format 'nativeFormat' is not supported." The vCenter Server returns this error in the following situations:

1. Create a VAAI NAS native linked clone virtual machine that uses `nativeFormat` for the delta disk format (`VirtualMachineCloneSpec.location.disk.diskBackingInfo.deltaDiskFormat`).
2. Create a second-level clone from the first clone, such that the disk for the second-level clone is on a different datastore. If you do not specify `seSparseFormat` or `redoLogFormat` for the second-level clone delta disk format, the clone operation will fail.

If you are using a virtual machine clone with a native delta disk format, you must specify either `seSparseFormat` or `redoLogFormat` for any clones that you create from the original native clone when the second-level clone is on a different datastore.

- **The `fault` parameter is required when setting the `SetTaskState` method.**

The *vSphere Web Services API Reference* does not document the following limitation with regard to the `SetTaskState` method:

If you specify an error state when you call the `SetTaskState` method, you must also specify the `fault` parameter in the calling sequence.

- **The `seSparse` virtual disk type is intended for internal use only.**

The documentation does not specify internal use only for `VirtualDiskType.seSparse`.

Resolved Issues

- None

Copyright © Broadcom