

Veritas™ Dynamic Multipathing readme

Solaris

5.1 RP1 DMP



Early release issues and notes

This document includes the following topics:

- [Early release issues and notes](#)

Early release issues and notes

In this release, Veritas Dynamic Multipathing has the following limitations when DMP is used with ZFS pools.

Device naming limitations

In order to use DMP for OS native device support, the device naming scheme must be enclosure-based naming (EBN) and the persistence attribute must be on.

The naming scheme cannot be changed after the feature is enabled.

Disk access name limitations

The following limitations apply to disk access names:

- You cannot change the DMP name of a device when native device support is enabled. If you must change the DMP name, you must disable the feature, update the name and again enable the feature.
- You also cannot change the name for an enclosure which has at least one DMP device that is using native device support, because that would require the DMP device name to change.
- When native multipathing is enabled, you cannot assign device names manually using the `vxddladm assign names` command.

Excluding arrays or paths from DMP

If an array includes at least one DMP device that is using native device support, you cannot exclude the array from DMP using the `vxddladm excludearray` command.

If a DMP device has the native device support enabled, you cannot exclude all of the paths from VxVM and DMP. You can exclude paths from VxVM as long as at least one path remains for the DMP device.

Third party driver and foreign device support

When DMP devices have OS native device support enabled, do not add the paths of the DMP device to VxVM using the under the foreign disk category.

If a DMP device has a valid VxVM label, you cannot enable native device support on that device.

You cannot enable native device support on a DMP device if the device is in co-existence with any Third Party Multipathing Driver (TPD) meta-device.

The Symantec testing teams have not yet certified the procedures to migrate ZFS pools over HDLM.

Solaris ZFS pools (zpool) operations take a long time

DMP support for native devices is enabled with the command `vxddmpadm settune dmp_native_support=on`. If this command is used on a system with existing ZFS pools, the ZFS pools (zpool) are moved onto DMP devices. This operation may take a long time depending on the time taken by the ZFS commands that are executed internally.

The following operations may take a long time:

- Importing a zpool can take a long time if you do not specify the scan directory to import the zpool. This is a Solaris ZFS limitation.
- The zpool operations take a long time due to passive not ready path. Currently, this issue is only seen with the STK 6540 array.
- If the zpool is imported on another host, the import fails. The import can take a long time before it fails due to the ZFS command that returns the error.

Using DMP with zpool on a simple disk

On Solaris, when a zpool is created on a whole (simple) disk (device name not containing any slice), ZFS changes the label on the disk from SMI to EFI label. The DMP device shows in the error state until you update the DMP data structures.

To update the DMP data structures with this change, run the `vxdisk scandisks` command.

Using DMP devices as hot spares

On Solaris, for an existing zpool with hot spares, the procedure to use DMP devices as the hot spares is as follows:

To use DMP devices as hot spares for an existing zpool

- 1 Remove the OS devices designated as spares from the zpool.

```
# zpool remove pool device
```

- 2 Add the corresponding DMP devices as spares.

```
# zpool add pool spare device
```

Documentation errata in vxdmadm manual page

The `vxdmadm` manual page is missing the description of the tunable 'dmp_native_support'.

