

Tru64 UNIX

Preparing TruCluster Version 1.5 or 1.6 Hardware and File Systems for an Upgrade to TruCluster Server Version 5.1

March 2001

This Best Practice describes how to perform advanced hardware and file system planning so that you can prepare your current TruCluster Available Server or Production Server Version 1.5 or 1.6 cluster for an upgrade to TruCluster Server Version 5.1.

Contents

Preparing TruCluster Version 1.5 or Version 1.6 Hardware and File Systems for an Upgrade to TruCluster Server Version 5.1

Preparing TruCluster Version 1.5 or Version 1.6 Hardware and File Systems for an Upgrade to TruCluster Server Version 5.1

See the Tru64 UNIX Best Practices Web page for more information about Best Practices documentation.

Is This Best Practice Right for You?

Not all Best Practices apply to all configurations, so you must be sure that it is appropriate for your system and circumstances. To use this Best Practice, you must meet the requirements described in the following table:

Requirement	Description
Operating System	DIGITAL UNIX Version 4.0D or later plus patches (see the Tru64 UNIX Software Patch Kits Web site)
System Configuration	TruCluster Available Server Software or Production Server Software Version 1.5 or 1.6
Impact on Availability	Depends on the current hardware and file system configuration
Skill Level	This Best Practice is intended for experienced cluster administrators.

If you do not meet these requirements, see *Alternative Practices* for information.

Before You Begin

Before you apply this Best Practice for preparing hardware and file systems for upgrading to TruCluster Server Version 5.1, you must:

- Understand your current cluster configuration
- Understand your upgrade options
- Know where you can find more information

Understanding Your Current and Future Cluster Configurations

Most of the preparation for an upgrade involves understanding the changes introduced in Tru64 UNIX Version 5.0 and how you want to configure your cluster for TruCluster Server Version 5.1. For example, you need to understand how the changes in device naming and file system namespace introduced in Tru64 UNIX Version 5.0 will impact your Version 1.5 or 1.6 cluster and its applications as you move toward Version 5.1. For more information about TruCluster Server Version 5.1 features, see the TruCluster Server Technical Overview manual.

Additionally, you need to relate Available Server Environment (ASE) concepts to the features and requirements of Version 5.1. If you understand your current cluster configuration and how you can obtain similar or improved functionality in TruCluster Server Version 5.1, you can design an upgrade approach that fits your needs. For more information about how to migrate ASE services to TruCluster Server Version 5.1, see the TruCluster Server Highly Available Applications manual.

Understanding Your Upgrade Options

There are several ways to approach upgrading your current cluster to TruCluster Server Version 5.1. You need to understand the following upgrade paths before making a decision:

- Option 1: New systems and new storage
- Option 2: New systems and existing storage
- Option 3: Existing systems and existing storage

See the TruCluster Server Software Installation manual for the benefits and issues associated with each approach.

Option 1: New Systems and New Storage

Build a separate cluster with new systems and new storage hardware. When the new cluster is fully configured and tested, use backup and restore mechanisms to migrate data from the old cluster to the new cluster.

Option 2: New Systems and Existing Storage

Build a separate cluster with new systems and enough new storage to install the cluster and test applications. When the new cluster is fully configured and tested, physically move the existing storage devices from the old cluster to the new cluster.

Option 3: Existing Systems and Existing Storage

Upgrade using existing hardware. Remove and disconnect one member from the old cluster, install and configure Tru64 UNIX Version 5.1, shut down the old cluster, connect storage devices to the Tru64 UNIX Version 5.1 system, configure storage, create a single-member cluster, and then add the other systems to the new cluster.

Where to Find More Information

See the TruCluster Server Software Installation manual for detailed guidelines about upgrading TruCluster Available Server Software or Production Server Software Version 1.5 or 1.6 to TruCluster Server Version 5.1.

Applying the Best Practice

Before you proceed, be sure to follow the recommendations in *Before You Begin*.

The following hardware and file-system requirements are the most common that can affect an upgrade from TruCluster Available Server Software or Production Server Software Version 1.5 or 1.6 (current cluster) to TruCluster Server Version 5.1 (new cluster):

- The current cluster must have required hardware. See *Step 1: The Current Cluster Must Have Required Hardware (Option 3)*.
- The current cluster should have a symmetrical shared storage configuration. See *Step 2: The Current Cluster Should Have a Symmetrical Shared Storage Configuration (Option 2 and Option 3)*.
- The new cluster must meet minimum disk requirements. See *Step 3: The New Cluster Must Meet Minimum Disk Requirements (All Options)*.
- RAID controllers must be used to mirror file systems in TruCluster Server Version 5.1. See *Step 4: Use RAID Controllers to Mirror the root (/) File System (All Options)*.
- UNIX File System (UFS) file systems must be migrated to Advanced File System (AdvFS) file systems. See *Step 5: Migrate UFS File Systems to AdvFS File Systems (All Options)*.

These requirements are explained in the following sections in more detail.

Step 1: The Current Cluster Must Have Required Hardware (Option 3)

The current cluster systems must be peripheral components interconnect (PCI) systems that are supported by TruCluster Server Version 5.1. The

systems must be PCI-based because TruCluster Server Version 5.1 requires Memory Channel as a cluster interconnect, and Memory Channel adapters use PCI slots.

Note

The DEC 3000, DEC 7000, DEC 10000, AlphaServer 300, AlphaServer 400, and AlphaServer 1000 are not supported in a TruCluster Server Version 5.1 cluster. See the TruCluster Server Software Product Description for a list of systems supported by TruCluster Server Version 5.1.

TruCluster Server Version 5.1 requires at least one PCI Memory Channel adapter on each system. If your current cluster consists of two systems, a Memory Channel hub is optional; the Memory Channel adapters can be connected with a cable. If more than two systems are in your current cluster, a Memory Channel hub is required.

Note

The DS20, ES40, GS80, GS160, and GS320 systems support only Memory Channel 2 hardware in a cluster interconnect. See the TruCluster Server Hardware Configuration manual for Memory Channel restrictions.

If the new cluster will require additional storage hardware, add this storage hardware to the current cluster *before* beginning the upgrade. Make sure that the new cluster has shared storage for the following file systems and disks.

Note

These are recommended minimum disk space requirements. We recommend that you exceed the minimum disk requirements for the shared root (/) file system and member boot disk root and swap areas on member boot disks by 25 percent to allow for future growth.

- Clusterwide file systems:
 - Cluster root (/) — 200 MB
 - Cluster /usr — 1000 MB
 - Cluster /var — 1000 MB

- Member boot disks:
 - Member boot disk root — 128 MB
 - Member boot disk swap — 128 MB
 - Member boot disk CNX cluster status — exactly 1 MB

- Optionally, a quorum disk (which is highly recommended for two-member clusters) — exactly 1 MB.

See the TruCluster Server Software Installation manual for more information about recommended disk space.

Tru64 UNIX Version 5.1 supports SCSI IDs 0-15. If you do not have SCSI IDs available before starting the upgrade, you can add additional storage after you boot Tru64 UNIX Version 5.1 and finish mapping existing device names to new device names. All storage must be attached and visible to the Tru64 UNIX system before creating a cluster.

The upgrade to TruCluster Server Version 5.1 is a good time to rearrange existing hardware, take advantage of new features such as multipathing, or move local storage to shared storage, if possible. Use this opportunity to improve your hardware configuration.

Step 2: The Current Cluster Should Have a Symmetrical Shared Storage Configuration (Option 2 and Option 3)

The current cluster should have a symmetrical shared storage configuration so that each shared device is known to all cluster members by the same special file name. For example, on all members, `rz17c` refers to the same physical device. The Option 2 and Option 3 upgrade procedures use scripts specifically designed to facilitate the migration of storage from the old cluster (`rz*` style device names) to the new cluster (`dsk*` style device names).

Supported Available Server configurations intrinsically use symmetric shared storage. Production Server clusters must be properly configured and contain only one Available Server Environment (ASE).

The automated scripts in the TCRMIGRATE510 subset of the TruCluster Server Version 5.1 kit depend on the current ASE environment having unique AdvFS domain names, LSM volumes, and device special file names for all shared storage. See the TruCluster Server Software Installation manual for more information about using the automated upgrade scripts in TruCluster Server Version 5.1.

Creating a configuration map of the current ASE can help you determine whether you can use the upgrade scripts to map storage during an upgrade. Use the `cluster_map_create cluster_name-full` command to create a configuration map of the current environment and the cluster monitor (`cmon`) to display and print the map. See the TruCluster Production Server Software and TruCluster Available Server Software Version 1.5 or Version 1.6 Administration manual for more information on the `cluster_map_create` and `cmon` commands.

If you cannot make your current storage configuration symmetrical, see *Alternative Practices*.

Step 3: The New Cluster Must Meet Minimum Disk Requirements (All Options)

The minimum recommended disk configuration for a TruCluster Server Version 5.1 cluster is as follows:

- One or more private disks on the system that will become the first cluster member *or* one or more disks on a shared bus that the system can access to hold the Tru64 UNIX operating system.
- One or more disks on a shared SCSI bus to hold the clusterwide `/`, `/usr`, and `/var` AdvFS file systems.
- One boot disk per member, normally on a shared SCSI bus.
- Optionally, one disk on a shared SCSI bus to act as the quorum disk.

For example, for a two-member cluster, you minimally reserve five disks (one for the private Tru64 UNIX operating system, one for the clusterwide file system, two for member boot disks, and one for the quorum disk). See the TruCluster Server Software Installation manual and the TruCluster Server Hardware Configuration manual for more information about disk requirements.

Step 4: Use RAID Controllers to Mirror the `root (/)` File System (All Options)

The `root()`, `/usr`, and `/var` file systems must be highly available in a cluster, so you should mirror these file systems. Because LSM is not supported for the `root(/)` file system (or member boot partitions, the quorum disk, or swap partitions) in TruCluster Server Version 5.1, you need to use RAID controllers (for example, an HSZ70) to mirror the `root (/)` file system. You can use either RAID controllers or LSM (software RAID) to mirror `/usr` and `/var`. Typically, using RAID controllers is more convenient to mirror all three of the file systems.

See the TruCluster Server Hardware Configuration manual for information about configuring storage in a TruCluster Server Version 5.1 cluster.

Step 5: Migrate UFS File Systems to AdvFS File Systems (All Options)

TruCluster Server Version 5.1 clusters use AdvFS file systems; UFS is supported as a read-only file system. If you use UFS file systems on the current cluster, we recommend that you convert those file systems to AdvFS before beginning the upgrade. How you perform data migration depends on whether you have enough storage in your current configuration to create AdvFS domains on new storage, or whether you must perform backups and restores to reuse current storage.

You can also convert to the Tru64 UNIX Version 5.1 AdvFS on-disk format (Version 4) after the upgrade to TruCluster Server is complete. The AdvFS on-disk format provides performance enhancements that are most noticeable when dealing with directories that contain thousands of files; for example, web servers. Whether you convert to the Version 4 format, and when you convert are up to you. Tru64 UNIX Version 5.1 recognizes both the Version 3 and the Version 4 formats. You can use the `mkfdmn -v3` option to create old-style AdvFS domains on the new cluster.

If you plan to reuse current storage, converting from UFS to AdvFS on the current cluster allows you to convert the file system at your convenience. If you do not convert on the current cluster, you must do the conversion as part of the upgrade, which increases your down time. You may also have a difficult time reverting to the old cluster if you encounter problems during the upgrade.

If you want to take advantage of the AdvFS Version 4 format on the new cluster, you have two UFS conversion choices:

- Convert to AdvFS before migrating, and then convert again to the AdvFS Version 4 format after migrating. You perform two conversions, but the data is available for read/write access as soon as it is migrated to the cluster.
- Convert to AdvFS after migrating. You perform only one conversion, but the data is available read-only until you perform the conversion.

See Appendix C of the Tru64 UNIX AdvFS Administration manual for instructions on how to convert UFS file systems to AdvFS.

Verifying Success

How you upgrade your Version 1.5 or Version 1.6 cluster to TruCluster Server Version 5.1 depends on several considerations. Following this Best

Practice to prepare hardware and file systems positions your current cluster for a successful upgrade.

If you were not able to successfully prepare hardware and file systems for an upgrade to TruCluster Server Version 5.1, see *Troubleshooting* for information about identifying and solving problems. See the TruCluster Server Software Installation manual for instructions on upgrading a cluster from Version 1.5 or Version 1.6 to TruCluster Server Version 5.1.

Troubleshooting

If you determine that the Best Practice was not successful, as described in *Verifying Success*, use the following table to identify and solve problems:

Problem	Possible Solutions
Do not have sufficient hardware	Make sure you have PCI-based systems and that you have available slots for at least one Memory Channel adapter.
Current cluster does not have symmetrical storage	If your current storage configuration is not symmetrical and you cannot make it so, upgrade to TruCluster Server Version 5.1 using Option 1.
Do not have enough disks for the new cluster	A two-member cluster requires a minimum of five disks. You can add additional disks as needed.
Do not currently use LSM to mirror root (/)	Use hardware RAID to provide the same no-single-point-of-failure (NSPOF) configuration in TruCluster Server Version 5.1.
Currently use UFS file systems	TruCluster Server Version 5.1 clusters use AdvFS file systems. Convert UFS file systems to AdvFS.

Alternative Practices

If you cannot address all requirements in this Best Practice, use Option 1 to upgrade to TruCluster Server Version 5.1. You can also use Option 2 if you have or can obtain the appropriate storage hardware.

With Option 1 you create a separate cluster with all new system hardware and storage. Option 1 lets you configure, test, and tune the new cluster without carrying over any of the limitations of the current hardware configuration, or any restrictions imposed by the Version 4.* operating system or the Version 1.5 or 1.6 TruCluster Software products. For example, you could build a new no-single-point-of-failure (NSPOF) cluster using HSG80 controllers and Fibre Channel.

Using a separate, but parallel, cluster with new systems and new storage has the lowest risk because no hardware is shared between the old and new clusters and no storage is migrated from the old cluster to the new cluster. In addition, you can extensively test applications on the new cluster while still serving clients from the old cluster.

With Option 2 you create a separate cluster with some new storage, but move storage hardware from the ASE to the new cluster. This approach provides a degree of separation in that you configure new systems and perform application testing without affecting the ASE.

See the TruCluster Server Hardware Configuration manual for information about how to design the hardware configuration for the new cluster. See the TruCluster Server Software Installation manual for detailed guidelines on upgrading your existing cluster to TruCluster Server Version 5.1.

Comments and Questions

We value your comments and questions on the information in this document. Please mail your comments to us at this address:

`best_practices@zk3.dec.com`

Legal Notice