



**Implementing TDMF UNIX
technology for simple, effective
and nondisruptive data migrations.**

A technical implementation guide

Contents

2 Introduction
2 TDMF UNIX capabilities
6 TDMF architecture
8 TDMF UNIX phases
12 TDMF performance considerations
15 Tips for conducting large migrations
18 Comparison with volume manager mirroring
19 Other Softek migration solutions
19 Platform support
20 Summary

Introduction

Softek™ TDMF™ UNIX® technology is host-based software that moves data at the block level between logical volumes (file system managed or raw volumes) without interrupting the applications that read/write to those volumes. Transparent Data Migration Facility (TDMF) software works with any volume manager, file system or database raw volume.

This software is ideal for the following types of local data migrations:

- *Technology refresh, especially in heterogeneous or high-availability environments*
- *DAS to SAN migrations*
- *Implementing tiered storage*
- *Consolidating storage locally*
- *Dynamic data movement to fix performance problems (e.g., moving volumes from storage hot spots)*

TDMF UNIX capabilities

Simple, easy installation

The distinct value of the TDMF UNIX solution is its simplicity—three components designed to execute flawlessly together. It is designed to be simple to install, simple to use and simple to maintain.

Highlights

With this technology, administrators can proactively prepare a system for its future migration before the details are defined.

Patent-pending nondisruptive switchover

TDMF's key differentiator is its ability to nondisruptively switch the source and target volumes, which moves the applications dynamically onto new storage. This switchover feature, which can occur automatically or on command, results in redirection of an application's READ function (e.g., from the original source to the target). This occurs without any disruption to the application.

TDMF continues to synchronously WRITE to source and target volumes, allowing switchback at any time. Switchback simply refers to the ability to repeatedly switch the source and target volumes. A switchback might be used, for example, to direct application READs back to the original source while tuning the configuration of the new storage volumes to improve performance.

Volumes that are no longer in applications' READ path can be taken offline without disruption.

Flexible target configuration

Once the TDMF software is installed and the source volumes are configured, target volumes can be identified at any time. This capability allows an administrator to proactively prepare a system for a future migration event, even when the details of that event are still undefined (e.g., new storage hardware has not yet been selected).

TDMF software supports the definition of up to eight target volumes for each source volume. This ability to create and synchronize multiple copies of data (known as "cloned data") is useful when supporting multiple test environments or variety of backup scenarios.

Highlights

When TDMF software is used, the number of volumes concurrently migrating data is configurable.

TDMF can pause the migration, reconfigure the parameters and resume the migration from where it left off.

Finally, target volume characteristics can be different from source volume characteristics. For example, the source may be a simple volume while the target may be a software RAID or striped configuration.

When TDMF software is used, the number of volumes concurrently migrating data is configurable. It is also important to note that logical unit numbers (LUNs) and storage layout can be changed during the data migration, since logical volumes are being moved.

Migration tuning parameters

To adjust performance impact on applications during a migration, TDMF provides the ability to control migration rate and CPU impact through performance tuning parameters, including:

- *Block size (determines the amount of data moved for each single read and write during the copy phase).*
- *Sleep interval (determines the pause interval between block movements).*
- *Concurrent copies in progress (determines the number of threads concurrently allocated to the copy phase of the migration). TDMF allows up to 50 concurrent volumes to be copied on a single copy command.*
- *Concurrent copy processes allowed on a single volume. TDMF allows up to four copy processes to be copying data on a single volume.*

Should business conditions require a change to these parameters, TDMF supports the ability to pause the migration, reconfigure the parameters and then resume the migration from the point it left off.

Highlights

***TDMF technology is designed to
work with any clustered system.***

Migration groups

TDMF supports the definition and concurrent migration of thousands of volumes. Because migration projects may involve several hundred volumes supporting a range of application types, TDMF provides the ability to logically group volumes for efficient operational control. Each group's migration parameters are independently configured and controlled to best suit the business applications supported.

Monitoring features

TDMF provides the ability to monitor the migration process from start to finish. Statistics that may be monitored include details such as estimated time to completion, elapsed time, time remaining, copy rate and percentage complete.

Cluster interoperability

TDMF is designed to work with any clustered system. Automation has been built for HP Service Guard clusters, IBM HACMP™ clusters, VERITAS clusters and Sun clusters. This cluster integration automates the majority of the processes needed to successfully complete data migrations in highly available cluster environments. It also reduces the complexity and risk in the most important production environments.

Highlights

The TDMF solution consists of three components: CLI, filter driver and configuration files.

TDMF architecture

TDMF consists of three components:

- *Command line interface (CLI): provides a user interface to issue commands and control migrations*
- *Filter driver: captures application updates during the data migration*
- *Configuration files: contains system and migration configuration details referenced by the TDMF software.*

The TDMF software is transparent to the application(s) running on the application server. Since the TDMF driver is located between the file systems and the volume manager, there are no driver, microcode or multipathing dependencies.

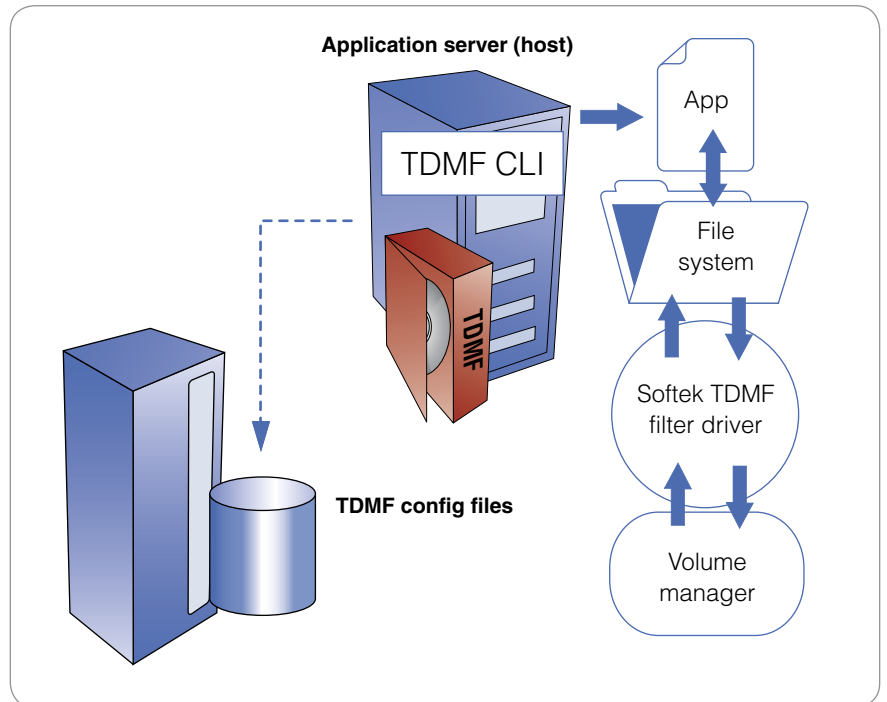


Figure 1
TDMF architecture

Highlights

TDMF UNIX technology uses a virtual migration device to support data integrity during dynamic swap and to capture the application updates during the migration.

TDMF virtual migration device

TDMF UNIX technology uses a virtual migration device to support data integrity during dynamic swap and to capture the application updates during the migration. This virtual migration device is a pseudo device that is created by using the CLI. One virtual device is created for each logical volume that will be migrated.

The directory structure for this device is /dev/mig. There are two entries created for each volume, a block device and special character device. The name of the virtual device is chosen by the administrator of the migration but the directory path must be /dev/mig. For example, if a migration volume is created with the name of oracledb, then the following virtual devices are created:

- /dev/mig/roracledb
- /dev/mig/oracledb

To activate TDMF, the file system entries for the volumes to be migrated must be changed in /etc/fstab to point to the TDMF virtual devices. Softek provides a script that automates this process, saving the current fstab and then making the appropriate changes, to put the TDMF software in the data path. To activate these changes the source file systems require an unmount and a remount, which normally requires a short application outage.

Highlights

The Common Console feature allows all TDMF migrations to be controlled centrally.

If the volumes that need to be migrated are raw volumes, then symbolic links will have to be used to ensure that the TDMF driver is in the data path. Normally raw devices are only migrations used by database applications such as Oracle or IBM Informix® technology. There are three ways to get the TDMF virtual device in the data path:

- *Update the database dictionary to point to the migration devices directly instead of the original volume.*
- *Create or update the existing symbolic links to point to the migration virtual devices.*
- *Rename the existing logical volumes and then create symbolic links that point to migration volumes which in turn point to the renamed volumes.*

TDMF UNIX phases

After the installation, an administrator can use the CLI to issue simple, easy commands to perform all steps of a migration process or can control the migration by using the Softek Common Console. This console is a centralized console that allows all TDMF migrations to be controlled centrally.

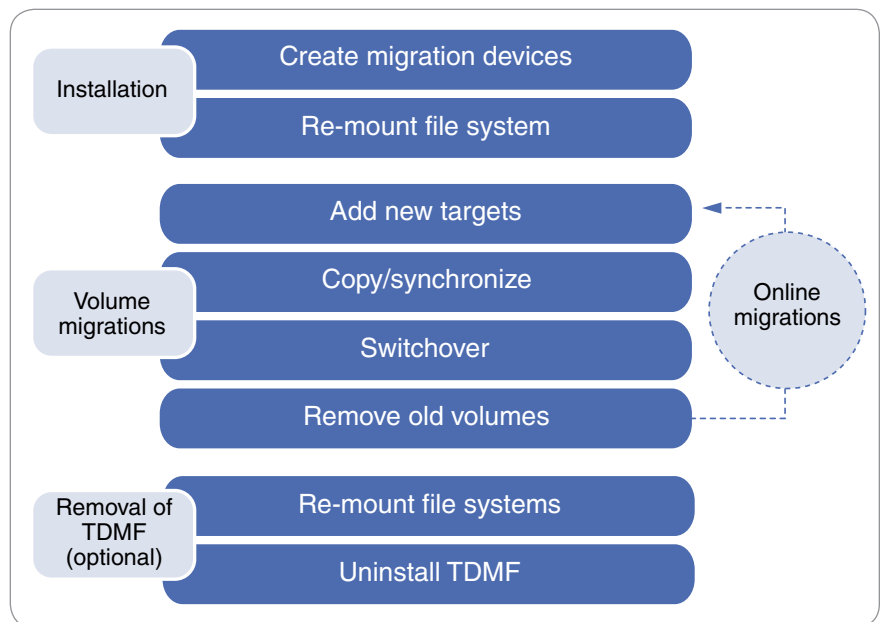


Figure 2
TDMF UNIX phases

Highlights

Once the migration volumes have been created, they can be grouped at any time for more efficient operational control.

Phase 1: migration volume creation

The first phase of a migration is to create the virtual migration devices for each source volume and update /etc/fstab as described in the previous section.

Once the migration volumes have been created, they can be grouped at any time for more efficient operational control.

Note that installing TDMF and creating/grouping the migration volumes can occur before target storage systems have even been decided upon. This lets an administrator proactively prepare a system for a future migration event. Often, this preparation is conducted during a normally scheduled downtime window or during off-peak hours (because of the application bounce) and can take less than 15 minutes.

Phase 2: addition of target volumes

Each migration volume requires at least one associated target volume of equal or greater size to start a migration copy process. Target volumes can be dynamically added to a migration configuration at any time, mapping a target volume to a migration volume (and its correlating source device).

Once the target volume has been added, the TDMF driver begins to send all updates synchronously to both the source and the target volumes. This ensures that, during the complete data migration progress, all updates that occur by the application will be applied to both the source and target volumes.

The initial status of a newly added target volume is SCRATCH, which means that the target volume is not usable by an application until a copy is completed (phase 3, data copy) and both the target and the source are synchronized.

Highlights

A migration can be paused or resumed at any time to modify parameters that increase or decrease the speed of the data copy process.

Phase 3: data copy

The copy phase is started on command by an administrator. The copy occurs asynchronously in the background while the applications are up and running. During the copy process, the status of target volume is COPYING.

The copy process can be paused and resumed during a migration. When paused and resumed, the copy operation begins from the point at which it was paused. It is also possible to start a migration from the beginning when paused if desired.

A migration can be paused/resumed at any time to modify parameters that increase or decrease the speed of the data copy process (called “throttling”). Throttling parameters provide an administrator with the ability to control subsystem input/output (I/O) to adjust the impact that a migration has on an application’s performance, including:

- *Blocksize of each read and write request (the greatest effect on speed)*
- *Sleep time used between each read/write copy request (10 ms default)*
- *Concurrent copy threads (four default)*
- *Concurrent copy processes used on a single volume (default 1).*

When the copy phase is completed, the migration volumes will automatically transition into the UPTODATE state. Once this transition occurs, then the migration volumes are eligible for the switchover phase. TDMF provides the ability to view the status of a migration at any time.

Highlights

The switchover process is designed to enable applications to remain online and 100 percent available.

Phase 4: switchover

Switchover is a patent-pending feature of the TDMF UNIX technology. This process is designed to enable the applications to remain online and 100 percent available. During the copy phase, all read requests are satisfied by the source disk device. When the switchover occurs, the source and target relationship is reversed. Following the switchover, all READ requests will be satisfied by the new volume (old target), while all updates (WRITES) are still applied to both the source and target volumes. The switchover is abstracted by the TDMF driver, which ensures full data integrity. All updates before and after the switchover are applied to both source and target volumes.

A switchover can occur by command when migration and target volumes are in an UPTODATE state (fully synchronized). Switchover can also occur automatically following a copy operation by using the -x option, which is automatic switchover after the copy is complete.

A switchover command can be repeated (called a “switchback”) in the event that a new storage volume is not performing as expected or configuration problems are encountered. Switchback is not possible once the old storage is removed.

If a switchover command is attempted while any of the migration volumes are not in the UPTODATE status, a switchover for that logical volume will not occur. Instead, an error message will be generated so that the migration administrator will be aware of the situation.

Highlights

After planning, the data copy phase is normally the longest migration phase.

Phase 5: old volume removal

When the switchover process is complete, the original source volumes can be removed from the migration pairs. When the old volumes are removed, all reads and writes are satisfied from the new storage device. By removing the original source volumes, the TDMF driver will no longer be writing to both the source and target volumes. The process of removing the original source volumes removes all TDMF allocation to these volumes. The volumes that have been removed from the migration can be varied offline and discontinued without taking an application outage.

This marks the completion of a successful data migration. If the TDMF software is to be uninstalled at a future time, then the fstab will need to be restored or changed to point to the new volumes.

TDMF performance considerations

Data copy speed

Next to planning, the data copy phase will normally be the longest elapsed time phase in the migration. Some recommendations for initial throttle settings:

- *Leave the `-s` parameter (sleep) value at the default for normal migrations*
- *Set the `-b` parameter (blocksize) to 8192, which is a good general setting for data migrations.*
- *Set the `-t` parameter (thread count) to 25, which gives the migration copy 25 concurrent copy threads.*
- *Set the `-n` (concurrent copy processes) parameter to 4 for volumes that are greater than 500 gigabytes in size.*

Highlights

Multiple copy commands can be run concurrently.

The limitation of concurrent copy threads is 50 per copy command. If this is not enough for the migration environment, multiple copy commands can be run concurrently, and there is no limitation in the number of concurrent copies that can be run. Of course, the storage subsystem eventually will be saturated, and no work other than the migration will occur.

Below is a sample table of results showing the impact of modifying throttle parameters during a migration. These measurements were taken on a small-scale fiber JBOD system. Results would scale dramatically if measured on an enterprise-level, high performance system.

-s 0 and -b param (K)	Actual MB/sec	Mig stat MB/sec
32 (New default)	6.8	7
64	12.1	12.5
128	19.8	20.5
256	29.4	30.1
512	23.6	24.3
1024	21.5	22.3
16384	18.4	19.4
32768	18.0	19.4
65536	18.4	19.4

-s 10 (default) and -b param (K)	Actual MB/sec	Mig stat MB/sec
32		292 KB/sec
512	4.0	4.1
1024	6.6	6.8
2048	9.7	10.1
4096	12.8	13.0
8192	14.6	15.0
16384	16.6	17.1
32768	16.8	17.7
65536	17.7	18.8

Highlights

TDMF software is designed to consume minimal CPU and memory resources.

CPU and memory consumption

TDMF software’s impact on the application and storage systems is very low. The majority of the resources consumed are I/O resources. The CPU consumption will normally be less than 5 percent of a single CPU for the software. The memory consumed by a single copy process will be less than 3 MB.

Below is an example of the resource consumption used during a migration and the recorded data transfer speed.

Environment setup	Migration test	Migration results
<ul style="list-style-type: none">• HP-UX 11.11• TDMF UNIX 2.0 software• Autopath Multipathing software• 2 Fiber HBAs• Source: HP EVA• Target: HP EVA• 100 one GB volumes	<ul style="list-style-type: none">• Copy using 50 threads on a single copy with 8192K blocksize	<ul style="list-style-type: none">• CPU consumption = 2.63 percent of a single CPU• Memory consumption = 2340K• Sustained IO rate = 140 MB/s

TDMF driver overhead

Softek has measured the overhead of TDMF driver when in synchronous mode (during the copy phase) with a 100 percent write workload. The results indicate an additional throughput and transfer speed overhead of about 25 percent.

To put these results into perspective: Most software mirroring drivers can add up to 100 percent overhead in the I/O latency. Given that a typical commercial workload has about a 20 percent write ratio, the overhead of the synchronous driver is actually greatly reduced in a live environment.

Highlights

With TDMF software, migration volumes can be added or deleted from groups without affecting the active migration.

Note that the described overhead is related to that period of time in which TDMF is copying data to synchronize the source and target systems, which is a few days on average. At all other times, TDMF adds extremely minimal overhead, especially when compared with traditional volume manager mirrors.

Grouping

TDMF has the ability to group volumes together for operational control. The migration volumes can be put into groups so that it is easier to control migrations with a single command(s) that is applied to all volumes in the group. All migration volumes will always be in the group of “all” so that commands can be entered with the scope of all or -a.

The grouping is decided upon by the migration administrator and can be changed dynamically. Migration volumes can be added or deleted from groups without affecting the active migration. It is recommended that the grouping follows some conventions such as by volume group name or by application or by array that is being migrated.

Tips for conducting large migrations

When multiple application requirements exist or when large servers are migrated, it is recommended to control the migration through the use of groups and dynamic switchover features.

As mentioned earlier, TDMF has the ability to group volumes together for operational control. Grouped migration volumes are easier to control; a single command affects all volumes in the group. All migration volumes will always be in the group of “all” so that commands can be entered with the scope of all or -a.

Highlights

The synchronous driver overhead is minimized because the dynamic switchover feature works by group, rather than by moving all volumes simultaneously.

The grouping is decided upon by the migration administrator and can be changed dynamically. Migration volumes can be added or deleted from groups without affecting the active migration. It is recommended that grouping follow some convention, such as by volume group name, application, or by array that is being migrated.

When a migration uses the dynamic switchover feature, you have the ability to start redirecting I/O application by application over to new hardware systems. Combined, these capabilities produce several favorable results:

- *The overhead of the synchronous driver is minimized because it is working by group, rather than bogging down the entire system by moving all volumes simultaneously.*
- *As volumes are migrated from the older system and application I/O is redirected to the new system, active I/O to the volumes on the original source is reduced, which can improve application performance through load balancing.*
- *As the migration proceeds and I/O is redirected volume by volume, applications can begin taking advantage of performance enhancements on the newer storage system for immediate service level improvements.*

Migration involving multiple applications

The following is an example of a migration where a storage hardware lease is expiring and data is being moved to a new higher-performing storage system. The storage system has 90 logical volumes defined, with 40 of them supporting a mission-critical, highly available Oracle application and 50 volumes supporting a data warehouse with predictable access patterns. Variable access exists during the week, but no access is available on weekends.

Highlights

The administrator groups the 40 Oracle volumes into Source Group A and the 50 warehouse volumes into Source Group B.

- *Source Group A: Migration is configured to occur nondisruptively during the week (i.e., running in the background) with synchronization expected in three business days. After peak business hours, the administrator accelerates migration speed slightly before going home; the adjustments affect all volumes in the group. In the morning, the administrator throttles back the migration of all volumes by modifying a single group parameter. By Friday, the target is synchronized with the source and a switchover occurs automatically.*
- *Source Group B: This group is configured to migrate as swiftly and rapidly as possible during weekend hours when the warehouse is rarely accessed and when the Oracle application from Group A is scheduled to be offline for two four-hour maintenance windows. This configuration maximizes the speed and IO usage of the machine while the Oracle application is offline. On Saturday, the administrator starts the copy operation for Group B volumes during the first four-hour downtime window. With a single command, the administrator pauses the migration when the Oracle application is put back online. The next day, the administrator resumes the copy operation for this group with a single command; migration is complete before the second four-hour downtime window expires and a switchover occurs to the new hardware.*

With both groups fully migrated and both applications now reading from the target device, the original source device can be removed at any time.

The grouping feature allows users to independently control different workloads during a data migration.

Without the grouping feature, the administrator would have had to issue copy, pause and throttle-change pause commands for each of the 90 volumes. The grouping feature allows for independent control of different workloads during a data migration.

Highlights

Several key factors make TDMF software a stronger migration solution than volume manager mirroring.

Larger-scale migrations

A similar approach can be used for larger migration scenarios. For example, suppose the migration involved 2,000 volumes supporting five applications. One approach could be to group volumes by application and conduct migrations one application at a time, using dynamic switchover to maintain a performance balance.

If the migration involved 2,000 volumes supporting a single application, another approach might be to group the volumes in 10 groups of 200. As each group of 200 is moved, an automatic switchover lets the old source volumes be removed or repurposed right away. This also removes any overhead introduced by synchronously writing to both the source and target volumes.

Comparison with volume manager mirroring

Some of the reasons why using the TDMF software is a better migration solution than volume manager mirroring include the following:

- *Volume manager mirroring is not always included in the base volume manager and can add monetary cost and complexity to a migration project.*
- *Volume manager mirrors only mirror data within a single volume group. Often, this is not an option when the new LUNs do not fit into existing volume groups. TDMF allows the copying of data between volume groups.*
- *Volume manager mirroring normally only allows copying one volume at a time within a volume group and TDMF allows many volumes to be copied at a time within volume groups.*
- *Volume manager mirroring does not allow the enterprise to change the volume or volume group characteristics during the migration.*
- *There are volume group types such as RAID 5 that are not supported by volume mirroring. TDMF can copy RAID 5 volume groups to other RAID 5 volume groups or to simple volumes.*

Highlights

TDMF software is designed to provide a fast, simplified, optimized migration solution for local migrations, while TDMF IP provides a broader set of features ideal for moving data over distances.

- *Volume manager mirroring has no pacing or restart option built in. These limitations can hurt application performance during a migration and can extend the length of a migration project when problems occur*
- *Most enterprises have a mix of volume managers and operating environments. TDMF software provides a single, standard method to perform a data migration in these environments*

Other Softek migration solutions

The Softek TDMF software was designed and optimized specifically for local data migrations in which data will not be moved across a Transmission Control Protocol/Internet Protocol (TCP/IP) network. Softek provides TDMF IP for migrations involving movement over distances (known as “global migrations”). TDMF IP uses TCP/IP to deliver data over a local area network (LAN) or wide area network (WAN).

Both Softek TDMF and Softek TDMF IP can be used in the same environment to address different migration project requirements. TDMF provides a very fast, simplified, optimized migration solution for local migrations; TDMF IP provides a broader set of features ideal for moving data asynchronously over distances.

Platform support

TDMF UNIX is currently generally available on:

- *HP-UX versions 11.0, 11.11, 11.23 on the PA-RISC and Intel® Itanium® platforms*
- *Solaris versions 7 through 10 on SPARC platforms*
- *IBM AIX® 4.3.3, 5.1, 5.2, 5.3 on RISC platforms*



Summary

TDMF is optimized for local data migrations involving heterogeneous storage devices. It is designed to be lightweight, easy to install and easy to use. TDMF can scale as a solution for small to large scale data migration projects. This software can be controlled locally on each server via the CLI or centrally via the Common Console.

TDMF is designed to ensure that applications maintain maximum availability even as the underlying storage infrastructure is phased out and taken offline. TDMF was designed specifically to help customers, storage vendors and integrators successfully and quickly move to new storage technology.

For more information

For more information about implementing TDMF UNIX, visit:

ibm.com/services/storage

© Copyright IBM Corporation 2007

IBM Global Services
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
12-07
All Rights Reserve

IBM, the IBM logo, AIX, HACMP, Informix, Softek and TDMF are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product and service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.