

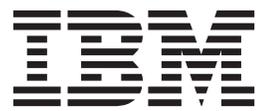
IBM FlashSystem

Integration Guide



IBM FlashSystem

Integration Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 21.

This edition applies to IBM FlashSystem and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2013, 2014.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

IBM FlashSystem Integration Guide . . . 1

Connectivity guidelines for improved performance of the IBM FlashSystem	1
Adding a second interface controller	1
Fibre Channel controller configuration guidelines	1
InfiniBand controller configuration guidelines	1
Host bus adapter configuration guidelines	2
Host channel adapter configuration guidelines	2
Cabling guidelines	2
Zoning guidelines	2
Logical unit number masking.	3
Sector size influence on performance	3
File alignment for best RAID performance	3
Storage modes for IBM FlashSystem 710 and IBM FlashSystem 810 systems	4
Maximum capacity storage mode	4
Active spare storage mode.	4
JBOF storage mode	4
Data acquisition storage mode	4
Microsoft Windows operating system configurations	5
Network topologies for Windows hosts	5
Sector size guidelines for Windows operating systems	5
Implementing 4 KB alignment for Windows 2000, Windows XP, Windows Server 2000 and Windows Server 2003	5
Windows Server 2003 multipathing.	6
Windows Server 2008 and Windows Server 2012 multipathing	6
Optimum disk command time out settings	6
Linux operating system configurations.	6
Network topology guidelines	6
Sector size guidelines for Linux	7
Aligning a partition using Linux.	7
Multipathing information for Linux.	7
Integrating InfiniBand controllers	8
Troubleshooting Linux configuration issues	8
Solaris operating system configurations	9
Network topology guidelines	9
Sector size restrictions for Solaris	9
Aligning the partition for Solaris.	9
Multipathing information for Solaris	10
Solaris 8 or Solaris 9 hosts	10
Solaris 10 hosts	10
Symantec Storage Foundation configurations	11
Sector size guidelines for Storage Foundation	11
Implementing alignment using Veritas Enterprise Administrator	11

Multipathing guidelines for Storage Foundation	11
Allowing for I/O recovery	11
VMWare ESX configurations.	11
Directly-attached Fibre Channel topology	11
Sector size guidelines for VMWare.	12
Alignment for VMWare	12
Multipathing for VMWare	12
HP-UX operating system configurations.	12
Directly attached Fibre Channel topology guidelines	12
Sector size restrictions	12
Alignment guidelines	12
Multipathing for HP-UX hosts	13
AIX operating system configurations	13
Directly attached Fibre Channel topology for AIX	13
Optimal logical unit number configurations for AIX	13
Sector size restrictions for AIX	13
Auto Contingent Allegiance support	13
Volume alignment	13
Implementing multipathing for AIX hosts	14
Resetting the host bus adapter and disk configuration	14
Setting the fast fail recovery flag for the host bus adapter	14
Rescanning for the storage system logical unit numbers	14
Checking the configuration	15
Configuring Active/Active path settings.	15
Oracle Automatic Storage Management	15
Microsoft SQL Server configurations	15
Disk and file system setup for SQL Server	16
Database layout	16
Moving an existing database to 4 KB sector storage	16
Migrating files to the storage systems	16
Detaching and attaching database files	16
Microsoft cluster services and Veritas.	16
Sybase configurations	17
Page size guidelines for Sybase.	17
Disk device configuration guidelines	17
Adding a mirror device	17
Sample multipath.conf file for Linux hosts	18
Verifying alignment using statistics tools	18

Notices 21

Trademarks	22
Terms and conditions	23

IBM FlashSystem™ Integration Guide

Learn about how to improve the performance of your storage system by applying specific connectivity, storage configuration, and alignment strategies for selected operating systems and applications.

Connectivity guidelines for improved performance of the IBM FlashSystem

You can configure the various network connections to improve the overall performance of the storage system.

Adding a second interface controller

You can install up to two controllers, each with two ports, in the storage system.

A second controller can increase the bandwidth and I/O operations per second (IOPS) performance, and add redundancy.

Fibre Channel controller configuration guidelines

You can improve reliability and performance by following specific network connection guidelines for Fibre Channel controllers.

Each dual-port 8 Gb Fibre Channel (FC) controller in the storage system provides bandwidth of close to 1,600 MBps. If two 8 Gb FC controllers are installed, there is an aggregate bandwidth of 3.2 GBps.

A single optical transceiver is the most likely component to experience failure on an FC controller. There are two hot-pluggable transceivers on an FC controller so a single transceiver does not cause the entire card to fail. However, both ports do share a common processor. To eliminate any FC failure point, use multipathing to span across two or more controllers before you use the second ports on the controllers. Most major operating systems incorporate a multipathing option. The input/output operations per second (IOPS) and bandwidth limits of the storage system increase with the number of controllers.

For optimal performance, connect the A ports of each FC controller before you use the B ports. Therefore, connect ports FC-1A and FC-2A before you connect ports FC-1B and FC-2B.

The storage system allows flexible presentation of LUNs through its FC controllers. A LUN can be presented and accessed through every port simultaneously. The loss of one controller does not affect the I/O performance on another controller. This configuration is commonly referred to as *Active/Active* or *Active/Active symmetric multipathing*.

InfiniBand controller configuration guidelines

Learn about network connections that add reliability and performance for InfiniBand controllers.

Each dual-port QDR InfiniBand controller provides bandwidth of close to 3 GBps. Installing two InfiniBand controllers in a storage system can provide an aggregate read-or-write bandwidth of 5 GBps across both InfiniBand controllers.

For optimal performance, connect the A ports of each controller before you use the B ports. Therefore, connect ports IB-1A and IB-2A before IB-1B and IB-2B.

The storage system allows flexible presentation of LUNs through its InfiniBand controllers. A LUN can be presented and accessed through every port simultaneously. The loss of one controller does not affect the

I/O performance on another controller. This configuration is commonly referred to Active/Active or Active/Active symmetric multipathing. Most major operating systems incorporate a multipathing option.

Host bus adapter configuration guidelines

Whenever possible, install dual-port host bus adapters (HBAs) in the server that is used with the storage system.

Dual-port and quad-port HBAs provide more ports for aggregating bandwidth. In an ideal case, multiple dual-port HBAs are installed in each server for redundancy.

Host channel adapter configuration guidelines

To improve bandwidth, install dual-port host channel adapters (HCAs) in the servers that are used with the storage system.

Dual-port HCAs provide more ports for aggregating bandwidth. In an ideal case, multiple dual-port HCAs are installed in each server for redundancy.

Quad Data Rate (QDR) HCAs maximize the bandwidth to the storage system. Most QDR HCAs use PCI Express 2.0 instead of version 1.0 to allow for the higher-rated bandwidth. Servers with PCI Express 2.0, or later, expansion slots provide the highest bandwidth.

The storage system is tested for interoperability against all major HCA vendors.

Cabling guidelines

Design cabling to provide resiliency across storage system ports, Fibre Channel (FC) or InfiniBand controllers, switches, and server host bus adapter (HBA) or host channel adapter (HCA) ports.

All high availability (HA) concepts of a dual-fabric setup apply to the storage system. When you connect the storage system to redundant fabrics, connect the A ports to one fabric and the B ports to the other fabric.

One key element to recognize when cabling a storage system is the use of available paths. The storage system is designed to deliver I/O throughput through all connected ports. Use all ports if the application can benefit from more bandwidth. The storage system supports up to two FC or InfiniBand controllers and supports four available ports. To take advantage of these ports on the storage system, there must be an equal number of server ports. Otherwise, there are underutilized ports on one side of the fabric.

Zoning guidelines

Implementing zoning when you connect the storage system to the network can improve performance and simplify upgrades.

In a switched fabric deployment, it is common to isolate one application's server-storage devices from other applications' server-storage devices. This practice prevents cross-traffic and helps ease maintenance. Therefore, employ zoning in all multiple-server environments.

Zoning is best deployed in a 1:1 port ratio with one server's host bus adapter (HBA) or host controller adapter (HCA) port for one storage system FC or InfiniBand port. A server HBA or HCA port that serves multiple FC or InfiniBand ports within the same storage system results in excessive paths because the server's single port is the limiting factor for performance.

Note: When you connect the storage system to a split-fabric architecture, ensure that each zoned server uses multiple controllers.

A suggested practice is to configure access for ports 1A and 2A, and then reverse the remaining ports 2B and 1B. The first server accesses ports on both controllers and does not depend on a single controller for availability.

Logical unit number masking

Logical unit number (LUN) masking is a utility that enables separate servers to access separate LUNs through a common controller port.

Although it is possible to use LUN masking in place of zoning, this function is better used as a supplement to zoning. If multiple servers access different LUNs through a common storage-system port, use LUN masking to present only the LUNs needed by the server.

Sector size influence on performance

If the application and file system supports it, a 4 KB sector size provides the best overall performance for the storage system.

The storage system includes the option of presenting each LUN with a unique sector size. These sizes range from a default of 512 bytes up to 4096 bytes (4 KB). Logical units that are configured with the 4 KB sector size provide optimal performance if the application and the underlying operating system are compatible with the advanced sizing. Most file systems can take advantage of 4 KB sectors, but the hosted application must also support the sector size.

Note: Format all file systems on the storage system at 4 KB, or at a multiple of 4 KB. For example, file systems that are formatted at 8 KB or 64 KB allocation sizes are satisfactory because they are a multiple of 4 KB.

Attention: Sector size changes are destructive to any existing file system on the LUN. You must choose the sector size before deploying the application to the storage systems.

File alignment for best RAID performance

File system alignment can improve performance for storage systems using a RAID storage mode.

File system alignment is a technique that matches file system I/O requests with important block boundaries in the physical storage system. Alignment is important in any system that implements a RAID layout. I/O requests that fall within the boundaries of a single stripe have better performance than an I/O request that affects multiple stripes. When an I/O request crosses the end point of one stripe and into another, the controller must then modify both stripes to maintain their consistency.

Unaligned accesses include those requests that start at an address that is not divisible by 4 KB, or are not a multiple of 4 KB in size. These unaligned accesses are serviced at much higher response times, and can also significantly reduce the performance of aligned accesses that were issued in parallel.

Your storage system provides 512-byte sector size support that greatly improves response times for I/O requests that cannot be forcibly aligned. However, alignment to 4 KB must be maintained whenever possible. When a 4 KB sector size is used on a LUN, there is no requirement to force alignment at the host or application.

For details about confirming that an application is submitting 4 KB aligned I/O requests, see “Verifying alignment using statistics tools” on page 18.

I/O alignment is visible only through write operations to the disk. If an application is seeing higher-than-expected response times for the write activity, contact support to help determine the optimal configuration for the application. For online support for these systems, see the IBM® Support portal.

Storage modes for IBM FlashSystem 710 and IBM FlashSystem 810 systems

Your storage system has storage mode options that increase flexibility and resiliency where appropriate.

Attention: Setting the storage mode alters the internal placement of data within the system and is not a data-in-place operation. You must back up all data before you change the storage mode.

Maximum capacity storage mode

The maximum capacity storage mode presents the highest usable capacity for creating logical unit numbers (LUNs).

Use this mode for deployments where the system is mirrored or preferred-read mirrored to other storage. In this mode, 240 Gb of capacity is presented for each flash module in an aggregate pool of capacity. LUNs are created across all available flash modules to achieve the highest performance.

Attention: If maintenance is performed that requires the replacement of a flashcard, all data must be restored from a backup after the maintenance is complete.

Active spare storage mode

Active spare storage mode improves system reliability.

The active spare storage mode reserves the last available card in the system as a spare for resiliency and serviceability. The most probable element of failure in the storage system is a flash chip. The Variable Stripe RAID (VSR) or RAID 5 layout of the memory within each flashcard provides protection against such an occurrence. If a flash chip fails, the active spare storage mode transfers data from the flashcard with a degraded chip onto the designated spare flashcard to restore full RAID protection.

The failover process is a data migration process that takes a few hours of nonblocking activity. The application can still operate under normal load while the data is copied to the replacement module. After the module completes its failover, schedule a maintenance window to replace the original module.

Note: Not all components are protected in this storage mode, so there is a risk of an unrecoverable failure. If a deployment requires no single point of failure, configure mirrored systems with the maximum capacity storage mode.

JBOF storage mode

Just a Bunch of Flash (JBOF) is analogous to the storage industry term JBOD (Just a Bunch of Disks).

The JBOF storage mode presents each flashcard as its own logical unit number (LUN). This mode is primarily meant to be used with host-level storage protection, such as a software RAID or Oracle Automatic Storage Manager (ASM).

Data acquisition storage mode

Data acquisition mode is appropriate for write-intensive applications.

The data acquisition storage mode is specific to specialized data acquisition applications and offers an increase in write performance. Use this mode only for write-only workloads that benefit from increased bandwidth.

Microsoft Windows operating system configurations

Learn how to improve performance of the storage system when configuring Microsoft Windows operating systems.

The storage system sees the most benefit from operating systems where multipathing and logical volumes are more fully supported. However, some applications are dependent on operating systems that are designed for workstations, and can still benefit from the storage system performance.

Network topologies for Windows hosts

Windows hosts have specific network topology requirements.

Arbitrated loop or point-to-point topology can be used on the Fibre Channel configuration for Windows hosts.

Sector size guidelines for Windows operating systems

Optimal sector support guidelines vary for systems running the Windows operating systems and applications.

Windows Server 2003 and Windows Server 2008 support the large 4 KB sectors and have good multipathing solutions. On the desktop series that run the Windows operating systems, multipathing is not available and 4 KB sector support is limited to Windows Vista and Windows 7.

Most Microsoft applications also support 4 KB sector support. SQL Server and Exchange support 4 KB sectors, but require a migration procedure when an existing data set is used.

If LUNs with a 512-byte sector size are used, alignment is critical on Windows Server 2003 and XP.

Implementing 4 KB alignment for Windows 2000, Windows XP, Windows Server 2000 and Windows Server 2003

Use this procedure to improve performance by establishing a 4 KB alignment on a Windows operating system.

About this task

Before Windows Vista and Windows Server 2008, systems running the Windows operating systems offset the partition by 63 sectors, or 31.5 KB.

To align to the preferred 4 KB sector size, implement the offset by using the diskpart.exe utility:

Procedure

1. To run the Windows diskpart.exe utility and view a list of available LUNs, enter the following command:
DISKPART> list disk
2. To select the LUN that holds the file system, enter the following command:
DISKPART> select disk <disk number>
3. To create a partition on the selected LUN, enter the following command:
DISKPART> create partition primary align=64
4. Use the Microsoft Management Console (MMC) or other method to assign a file system or drive letter (raw access) to the partition.

If you are installing the Windows XP or Windows Server 2003 operating system, create the partition on the LUN before the installation of the operating system. You can create the partition by using a *Linux Live* CD, or by presenting the LUN to another Windows host and disconnecting the drive after the partitioning is complete.

Windows Server 2003 multipathing

Configuring multipath I/O (MPIO) on Windows Server 2003 operating system can improve reliability.

Windows Server 2003 has a built-in multipath I/O (MPIO) driver that is provided by Microsoft. This driver is responsible for aggregating the links of storage systems and reporting the addition or removal of links to the kernel while online. To use this feature, a Device Specific Module (DSM) is provided to identify the storage system to the MPIO driver. All major storage vendors have support for the multipathing function and coexist safely because of the common driver in the Windows operating system. You can obtain a copy of this driver from support. For online support, see the IBM Support portal.

Note: EMC PowerPath data path management software must be upgraded to version 4.6, or later, before the storage system DSM is used.

Windows Server 2008 and Windows Server 2012 multipathing

Learn about configuring multipathing on servers by using Windows Server 2008 or Windows Server 2012 operating systems.

Windows Server operating system versions that begin with Windows Server 2008 no longer require a separate DSM. Instead, the multipath I/O (MPIO) function must be installed on the server. For more information, see the Microsoft TechNet website ([http://technet.microsoft.com/en-us/library/ee619752\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/ee619752(WS.10).aspx))

After you install the MPIO function, set the load balance policy on all storage system LUNs to **Lowest Queued Path**. All available paths to the LUNs are then used to aggregate bandwidth. The load balance policy is set through the Properties pane of each individual multipath disk device in Windows Device Manager.

Optimum disk command time out settings

Adjust the disk **TimeOutValue** parameter on the Windows operating system for more reliable multipath access.

Windows operating systems have a default disk command **TimeOutValue** of 60 seconds. If a SCSI command does not complete, the application waits 60 seconds before an I/O request is tried again. This behavior can create issues with most applications, so you must adjust the disk **TimeOutValue** in the registry key to a lower value. For more information about setting this value, see the Microsoft TechNet website.

- For a disk in a nonclustered configuration, set **TimeOutValue** to 10.
- For a disk in a clustered configuration, set **TimeOutValue** to 20.

Linux operating system configurations

Learn about improving storage system performance by using specific Linux configuration settings.

Most Linux distributions have the same optimum configurations.

Network topology guidelines

Linux hosts have specific network topology requirements.

You can use an arbitrated loop or point-to-point topology on Fibre Channel configurations for Linux hosts.

Sector size guidelines for Linux

Optimal sector support guidelines vary for Linux applications.

The Linux operating system supports large 4 KB sectors and also performs translation for most applications that still depend on 512-byte sectors. This means that applications that are not aware of sector sizes can use the 4 KB sector performance. The file system handles the small-block write operations that cannot be performed on 4 KB sectors with direct I/O.

For clustered applications where direct I/O is implemented, configure the application with 512-byte sector storage and alignment.

Aligning a partition using Linux

Use this procedure to improve performance by aligning a partition in the Linux operating system.

About this task

The Linux operating system defaults to a 63-sector offset. To align a partition in Linux using `fdisk`, complete the following steps:

Procedure

1. At the command prompt, enter `# fdisk /dev/mapper/ramsan`
2. To change the listing of the partition size to sectors, enter `u`.
3. To create a partition, enter `n`.
4. To create a primary partition, enter `p`.
5. To specify the partition number, enter `1`.
6. To set the base sector value, enter `128`.
7. Press Enter to use the default last sector value.
8. To write the changes to the partition table, enter `w`.

Results

The newly created partition now has an offset of 64 KB and works optimally with an aligned application.

If you are installing the Linux operating system on the storage system, create the partition scheme before the installation process. For most Linux distributions, this process requires starting to the text-based installer and switching consoles (Alt+F2) to the command prompt before you continue.

Multipathing information for Linux

You can use multipath I/O (MPIO) to improve performance of the Linux operating system.

Linux kernels of 2.6, and later, support multipathing through `device-mapper-multipath`. This package can coexist with other multipathing solutions if the other storage devices are excluded from `device-mapper`. For a template for the `multipath.conf` file, see “Sample `multipath.conf` file for Linux hosts” on page 18.

Because the storage system controllers provide true Active/Active I/O, the `rr_min_io` field in the `multipath.conf` file is set to 1. This results in a round-robin distribution of I/O activity across all available paths. If the I/O activity is more sequential in nature, you can increase the `rr_min_io` field by factors of 2 for a performance gain by using buffered I/O (non-DIRECT).

For 4 KB Sectors support, the `PathSelector` option in the `multipath.conf` file must be set to `tur` (test unit ready).

Integrating InfiniBand controllers

Learn about configuring the Linux operating system to work with InfiniBand controllers.

To integrate with InfiniBand technology, the storage system provides block storage by using the SCSI remote direct memory access (RDMA) Protocol (SRP).

The Linux operating system requires several software modules to connect to the storage system through InfiniBand technology and SRP. In particular, make sure you install the *srp* and *srptools* modules, and install drivers for the server's host channel adapter (HCA). Use the OpenFabrics Enterprise Distribution (OFED) package from <http://www.openfabrics.org> to install these modules, either individually or by using the Install All option.

For SRP to work with the storage system, make sure that the following settings in the `/etc/infiniband/openib.conf` file are set as follows:

```
# Load SRP module
SRP_LOAD=yes #
Enable SRP High Availability daemon
SRPHA_ENABLE=yes
SRP_DAEMON_ENABLE=yes
```

These settings cause the SRP and the SRP daemons to load automatically when the InfiniBand driver starts. The SRP daemon automatically discovers and connects to InfiniBand SRP disks.

Use the `SRPHA_ENABLE=yes` setting. This setting triggers the multipath daemon to create a multipath disk target when a new disk is detected.

InfiniBand technology also requires a Subnet Manager (SM). An existing InfiniBand network already has an SM. In many cases, an InfiniBand switch acts as the SM. If an SM is needed, install OpenSM, which is included with the OFED package, and start it on a single server in the network by entering the following command:

```
# /etc/init.d/opensmd start
```

This script opens an SM only on a single port. If multiple ports are directly connected to the storage system, a custom script is needed to start the SM on all ports.

Troubleshooting Linux configuration issues

Learn about troubleshooting Linux configuration issues.

- When you install OFED-X.X.X, an error occurs saying that the installer failed to build the ofa_kernel RPM Package Manager (originally Red Hat Package Manager).
 - The kernel that is used by the server might not be supported by OpenFabrics Enterprise Distribution (OFED). If the *Install All* option was chosen, try the *Customize option* in the OFED installation menu and select only the components that are needed. If this process does not work, try installing a different version of OFED.
- Loading the driver module fails.
 - The HCA used might not be supported by OFED, or the driver was not been installed correctly. Obtain the latest drivers for the host channel adapters (HCA) from the HCA vendor's website.
- When you try to install OFED, an error occurs such as `<module 1> is required to build <module 2>`.
 - This error means that some dependencies that are required by OFED are not installed on the server. You must install all of the required dependencies.
 - To search for the RPM needed (if the yum package-management tool is available) enter:

```
# yum provides <dependency_name>
```
 - To install the RPM, enter:

```
# yum install <dependency_rpm>
```

Note: If yum is not installed on the server, each dependency must be manually downloaded and installed.

-
- When you try to run the `srp_daemon` command, a message is displayed stating that an operation failed.
 - Make sure that the storage system is physically connected to the network and that all components are powered on.
 - Make sure that the correct cable is being used and that OpenSM is running. To confirm whether OpenSM is running, enter:

```
# /etc/init.d/opensmd status
```
- Loading the `ib_srp` module fails.
 - Check that OFED is installed correctly and that the necessary device drivers are also installed. If a custom OFED installation was performed, make sure that `ibutils` and all packages that are related to `srp` were selected.

Solaris operating system configurations

Learn about configuration guidelines for Solaris.

The Oracle Solaris operating system has some slight differences between x86 and SPARC support when the disks are partitioned. The MPxIO multipathing setups, however, are identical.

Network topology guidelines

Configure the settings of the Fibre Channel ports in the system for the network topology that is used with the storage system.

If directly attached to a server, configure the Fibre Channel (FC) ports in the storage system to an arbitrated loop topology. If configured through a switch, set the FC ports to point to point to properly negotiate with the switch.

Sector size restrictions for Solaris

You cannot implement a 4 KB sector size to improve performance if the host system is using the Solaris operating system.

The Solaris operating system does not support a 4 KB sector size, so all logical units presented to a Solaris host must be configured with a 512-byte sector size.

Aligning the partition for Solaris

Aligning the partition to a 4 KB boundary improves performance of the storage system.

About this task

Solaris SPARC aligns slices on logical unit numbers (LUNs) when the LUN is using an Oracle Sun label for a LUN smaller than 2 TB. If the 2 TB capacity is exceeded, the OS must use an Extensible Firmware Interface (EFI) label that causes slices that are not aligned to 4 KB boundaries.

EFI Disks default to 34-sector offsets that have the default partition table. To align the partition to a 4 KB boundary for optimal performance, complete the following steps:

Procedure

1. After formatting the disk, select the partition option.
2. Choose the All Free Hog partitioning base option.

```
partition> 0
Select partitioning base:
0. Current partition table (unnamed)
1. All Free Hog
Choose base (enter number) [0]? 1
```

3. Change the base sector to 40.
4. Press Enter to use the default values for the remaining options.

```
partition> 0
Enter partition id tag[usr]:
Enter partition permission flags[wm]:
Enter new starting Sector[34]: 40
Enter partition size[10066067388b, 10066067427e, 4915071mb, 4799gb, 4tb]:
Part Tag Flag First Sector Size Last Sector
0   usr  wm      40  4.69TB  10066067427
```

5. After the first sector is changed, save the configuration and then continue by using the newly created partition.

Multipathing information for Solaris

Learn about how to implement multipathing on Solaris hosts.

Solaris 8 or Solaris 9 hosts

The method of implementing multipathing depends on the host bus adapter (HBA) used in the server.

MPxIO is the built-in multipathing mechanism for Solaris and requires an HBA that is branded Oracle Sun. If you use an HBA that is not branded Oracle Sun, multipathing can be achieved by using Veritas Storage Foundation.

To enable MPxIO support, modify the `/kernel/drv/scsi_vhci.conf` file by using the following options:

```
load-balance="round-robin";
mpxio-disable="no";
device-type-scsi-options-list =
"IBM    FlashSystem", "symmetric-option";
symmetric-option = 0x1000000;
```

Important: There must be exactly five spaces between IBM and FlashSystem.

Solaris 10 hosts

The method of implementing multipathing depends on the host bus adapter (HBA) used in the server.

MPxIO is the built-in multipathing mechanism for Solaris and requires an HBA that is branded Oracle Sun. When you use an HBA that is not branded Oracle Sun, multipathing can be achieved by Veritas Storage Foundation.

To enable MPxIO support, modify the `/kernel/drv/scsi_vhci.conf` file with the following changes:

```
load-balance="round-robin";
device-type-scsi-options-list =
"IBM    FlashSystem", "symmetric-option";
symmetric-option = 0x1000000;
```

Important: There must be exactly five spaces between IBM and FlashSystem.

In addition, the `/kernel/drv/fp.conf` file requires the following change:

```
mpxio-disable="no";
```

Symantec Storage Foundation configurations

Learn about implementing Just a Bunch of Disks (JBOD) support for the storage system using Symantec Storage Foundation.

Symantec Storage Foundation requires that you run the following command to acknowledge multiple logical unit numbers (LUNs). You must run this command on the host:

```
# vxddladm addjbod pid="FlashSystem" length=10
```

After you enter this command, the Storage Administrator will list the LUNs appropriately.

Sector size guidelines for Storage Foundation

Some versions of Symantec Storage Foundation do not support the most efficient 4 KB sector size.

For versions through 5.1, logical units on the storage system must be configured with a 512-byte sector size.

Implementing alignment using Veritas Enterprise Administrator

Learn how to implement alignment when you allocate disks that use the Veritas Enterprise Administrator (VEA) interface.

When you allocate disks that use VEA, the default offset is unaligned to a 4 KB boundary.

To create an offset and to enforce track alignment on the disks before you introduce them into a volume, enter the following command:

```
# vxdisk set track align=enable offset=128 <disk>
```

Note: The `vxdisk` command applies only to Windows installations.

Multipathing guidelines for Storage Foundation

The storage system supports Active/Active multipathing through Veritas Dynamic Multi-Pathing.

Allowing for I/O recovery

The timeout threshold value affects the failover process.

To allow for failover in mirrored environments, lower the timeout threshold. To lower the timeout threshold, enter the following command:

```
# vxddmpadm settune dmp_failed_io_threshold=60
```

VMWare ESX configurations

Learn about improving performance in VMWare ESX environments.

VMWare ESX and VMWare ESXi are guest operating systems. Follow appropriate guest operating system guidelines through all levels of the application to achieve optimal performance.

Directly-attached Fibre Channel topology

Configure Fibre Channel controllers for arbitrated loop topology when directly attached to the ESX hosts.

Sector size guidelines for VMWare

ESX does not support a 4 KB sector size.

LUNs configured for 4 KB sectors cannot be addressed as passthrough devices.

Alignment for VMWare

Learn about aligning partitions on virtual machines (VMs).

VMWare aligns its datastores to 64 KB appropriately, but guest VMs must still align their own presentation of the storage. Before you continue the installation of a guest Linux operating system or a guest Windows Server 2003 operating system, partition the storage to the aligned accesses.

For more information, see “Linux operating system configurations” on page 6 and “Microsoft Windows operating system configurations” on page 5 regarding installing the operating system on aligned partitions.

Multipathing for VMWare

Learn about configuring multipathing for VMWare.

About this task

To configure round-robin multipathing, complete the following steps:

Procedure

1. In the vSphere client, select the **Configuration** tab.
2. In the **Devices** view, select each disk that you want to change the path selection on.
3. In the **Manage Paths** pane, change the Path Selection setting to **Round Robin (VMware)**.

HP-UX operating system configurations

Learn about improving performance with hosts that are using the Hewlett-Packard UniX (HP-UX) operating system.

Directly attached Fibre Channel topology guidelines

Learn about optimal network topologies for Fibre Channel controllers for Hewlett-Packard UniX (HP-UX) hosts.

Configure Fibre Channel (FC) controllers for an arbitrated loop topology. For HP-UX 11.11, to maintain consistent device naming across server restart operations, the FCs must be configured to a hard-loop ID assignment.

Sector size restrictions

The Hewlett-Packard UniX (HP-UX) operating system does not support 4 KB sector size.

All logical units must be configured with a 512-byte sector size.

Alignment guidelines

Learn about configuring the VERITAS File System (VxFS) for optimal alignment.

Set the VxFS to a block size of 4096, or higher, to maintain alignment.

To align the volumes to 4 KB boundaries, enter the following command:

```
# mkfs -F vxfs -o bsize=4096 <disk>
```

Multipathing for HP-UX hosts

Learn about implementing multipathing for hosts that are using the Hewlett-Packard UniX (HP-UX) operating system.

HP-UX v1 and v2 include Active/Passive multipathing software called PVLinks. The operating system reports multiple disk devices for a single logical unit (LU) with each device reflecting a Fibre Channel (FC) path to the disk.

To use PVLinks, add each disk and path into the same volume group under the Logical Volume Manager (LVM).

HP-UX v3 has automatic multipathing with a round-robin policy. No additional steps are required to use multipathing.

AIX operating system configurations

Learn about improving performance for hosts that use the AIX® operating system.

Directly attached Fibre Channel topology for AIX

Learn about optimal network topologies for Fibre Channel controllers that are connected to hosts that use the AIX operating system.

Configure the Fibre Channel controllers in the storage system to arbitrated loop topology when the controllers are directly attached to the AIX hosts.

Optimal logical unit number configurations for AIX

The number of logical unit numbers (LUNs) that you create on the storage system impacts overall performance.

Applications perform optimally if at least 32 LUNs are used in a volume group. If fewer volumes are required by an application, use the logical volume manager to map four logical volumes to 32 logical units.

Sector size restrictions for AIX

The AIX operating system does not support a 4 KB sector size.

All logical units must be configured with a 512-byte sector size.

Auto Contingent Allegiance support

You must enable Auto Contingent Allegiance (ACA) support on logical unit numbers (LUNs) that are connected to an AIX host by using the logical unit option in the storage system.

Some host systems require this ACA support to run multiple concurrent commands. After this option is changed, all interface ports that have access to the logical unit must be reset. In the web interface, select the port node in the tree and click **Reset Link**.

Volume alignment

The AIX operating system volumes align to 4 KB boundaries.

Implementing multipathing for AIX hosts

Learn how to implement multipathing for AIX hosts.

Important: For the latest updates for multipathing support on the AIX operating system, go to IBM Fix Central.

Resetting the host bus adapter and disk configuration

Learn how to reconfigure the host bus adapters (HBAs) to implement multipathing.

About this task

After you install the latest AIX updates for support of IBM FlashSystem™ storage, AIX must rescan for the LUNs to recognize them as devices that support multipathing. It is useful to start by reconfiguring the HBA and its attached disks.

Attention: If other disks are attached to any HBA devices, the following commands remove the configuration for those disks and the HBA. If you are attempting to save the current configuration, skip this task.

To reset the host bus adapter or disk configuration, complete the following steps:

Procedure

1. To determine the device names of the HBAs that the storage system is connected to, enter the following command:

```
lsdev -t efscsi
```
2. For each HBA device name, enter the following command to remove the HBA and the disk configuration that is associated with it:

```
rmdev -l <device name> -R
```
3. To determine whether any disks are already defined that must be removed before rescanning, enter the following command:

```
# lsdev -C -c disk
```
4. If any LUNs are already defined as Other FC SCSI Disk Drive, remove the old definitions. For each disk name, enter the following command:

```
rmdev -l <disk name> -d
```

Setting the fast fail recovery flag for the host bus adapter

You can set the fast fail recovery flag for the host bus adapter (HBA) to improve failover response.

For the multipath I/O (MPIO) driver to fail over to an available path in a timely manner after a path failure, set the `fast_fail` recovery flag for the HBA devices that the storage system is connected to.

At a command prompt, enter:

```
chdev -a fc_err_recov=fast_fail -l <device name>
```

where *device name* is the device name of the HBA that is connected to the system.

Rescanning for the storage system logical unit numbers

After the host system is configured to recognize that the storage device supports multipathing, you must rescan for the LUNs.

At a command prompt, enter:

```
cfgmgr -vl <device name>
```

where *<device name>* is the device name of the HBA connected to the system.

Checking the configuration

After you change the configuration to support multipathing, confirm that the configuration is working correctly.

About this task

To check the new configuration, complete the following steps:

Procedure

1. To ensure that the configuration was successful, enter `lsdev -C -c disk` at the command prompt to list all disks available to the system. All logical unit numbers (LUNs) should be shown as `MPIO IBM FlashSystem Disk`.
2. If there are missing disks, or extra disks, or the LUNs do not show as `MPIO IBM FlashSystem Disk`, check that the connections and the storage system configuration are correct. You may need to remove the configuration for the host bus adapters (HBAs) and complete the rescan again. For more information, see “Resetting the host bus adapter and disk configuration” on page 14.
3. Next, ensure that all the paths that are connected are visible. Enter `lspath` at the command prompt.
4. If paths are missing, check that the connections and the storage system configuration are correct. You must then remove the configuration for the HBA(s) and perform the rescan again. For more information, see “Resetting the host bus adapter and disk configuration” on page 14.

Configuring Active/Active path settings

Learn how to set up a storage device to actively use all paths to it.

Set the `algorithm` attribute to `round_robin` before you add the `hdisk` to any volume group, as shown in the following example:

```
chdev -l hdisk1 -a algorithm=round_robin
```

Oracle Automatic Storage Management

Learn about improving the performance of Oracle Automatic Storage Management (ASM).

Oracle ASM supports the use of 4 KB sectors, limited to the Linux operating system when presented through ASMLib. Oracle ASM is aligned to 4 KB sector boundaries by default and does not require 4 KB alignment to achieve optimal performance.

If a partitioned disk is presented to Oracle ASM, that partition must be aligned by following the alignment steps for the OS.

All Oracle table spaces, including User table spaces, Undo, and Temp, perform I/O activity at `db_block_size` multiples. If the database block size is 4 KB or higher for each table space, the I/O performance is optimal.

The Oracle Redo and Control files cannot be forcibly aligned. For these particular files, you can use the Log logical unit Number (LUN) feature of the storage system to accelerate the latency of file operations. A separate diskgroup that is presented from Log LUN units from the storage system is created for the Oracle instances of these specific files. For more information about the Log LUN feature, see your storage system user's guide.

Microsoft SQL Server configurations

Learn about improving the performance of Microsoft SQL Server.

Disk and file system setup for SQL Server

Learn about how to configure the file and disks for optimal performance of SQL Server.

SQL Server performs best with an underlying sector size on the storage of 4 KB. Format the New Technology File System (NTFS) at a 64 KB allocation unit size.

There is no noticeable performance advantage in separating database files, log files, and tempdb files. Separate these objects for easier management and performance monitoring.

Database layout

The database configuration can affect performance of the storage system.

Provision multiple database files (.ndf files) for each database to address concurrent I/O requests. The Microsoft recommendation is one database file per processor core. This practice also applies to the tempdb file.

Moving an existing database to 4 KB sector storage

Learn about effective methods of moving a database.

Migrating files to the storage systems

Use the **Copy Database** wizard in SQL Server to migrate files.

If an existing database created on a storage device with 512-byte sectors is being migrated to the storage system, use the **Copy Database** wizard in SQL Server.

Detaching and attaching database files

Learn how to implement a faster detach and attach workaround when you migrate the database to the storage system.

About this task

To reduce the time that is needed for the deployment, complete the following steps:

Procedure

1. Detach the database.
2. Move only the .mdf and .ndf files to the new storage.
3. Delete the original .ldf log file.
4. To attach the database files, open the Attach Database pane and select the mdf data file. By default, the list of files includes the original .ldf file as Not Found.
5. Remove this reference to the .ldf file so that the list includes only the .mdf and .ndf files.
6. Press OK.
7. To make the directory of the new log file point to the storage system, change the properties and restart the SQL Server.

Microsoft cluster services and Veritas

Learn about using Microsoft cluster services for highly available SQL database deployments on the storage system.

Use Microsoft cluster services to mirror the storage system if it is an SQL database deployment that is highly available. This configuration supports advanced deployments, including preferred-read mirroring between the storage system and a slower traditional hard disk drive system. The configuration also supports network-based replication.

Veritas Storage Foundation is the supported cluster volume manager for cluster services and the storage system. For more information, see “Symantec Storage Foundation configurations” on page 11.

Sybase configurations

Learn how to improve the performance of a Sybase database.

Page size guidelines for Sybase

Learn about page sizes for Sybase.

The default page size for Sybase is 4 KB. The page size can be 2 KB, 4 KB, 8 KB, or 16 KB. Most write operations occur at eight-page collections called *extents*. For a 2 KB page size, an extent is 16 KB. Page size is universally set on the Master Database.

Disk device configuration guidelines

Learn about configuration details for raw disk devices and file-system disk devices for Sybase.

Sybase works with raw disk devices and file-system disk devices.

For raw devices, configure the `/dev/mpath/<alias>` path for multipath targets, because the `/dev/dm-X` value varies and the `/dev/mapper/<alias>` path does not function.

File system devices can have a setting of **DSYNC**, **DirectIO**, or **CachedIO**.

- Use the **DSYNC** setting for devices when the page size is 2 KB.
- Use the **DirectIO** setting for devices when the page size is 4 KB or higher.
- Use the **CachedIO** setting for hard disk drives, not for flash storage. The setting is optimized for hard disk drives, not for flash storage, and therefore must not be used for the storage system.

You can use System Procedure `sp_logiosize` to modify the write size of the log file to a 4 KB value for better performance and space efficiency.

Adding a mirror device

Learn about adding a mirror device in Sybase, to provide redundancy if a database device fails.

To add another storage system as a mirror device, you must first enable disk mirroring and restart the database.

To enable disk mirroring and restart the database, enter the following commands:

```
exec sp_configure 'disable disk mirroring', 0
go

<db restart>

USE master
go
disk mirror name='fs1', mirror='/dev/mpath/fs2', writes =
'noserial'
go
```

Note: The `/dev/mpath/rs2` path cannot be a database device that is already defined.

After the device is mirrored, the database can read only from the primary copy of the mirror. This procedure can also be used for preferred-read mirroring. Preferred-read mirroring takes advantage of the storage system's speed for read-intensive data.

Sample multipath.conf file for Linux hosts

Use this sample multipath.conf file to implement multipathing for Linux hosts.

Note: This sample multipath.conf is valid for the following Linux distributions:

- Red Hat Enterprise Linux (RHEL) 6.x and 5.x
- Novell SUSE Linux Enterprise Server (SLES) versions 10 and 11

Note: For IBM System z[®], the following options are supported:

- IBM z/VM[®] with supported Linux for System z Guest OSes
- SLES 11 SP2 (native Linux for System z)
- RHEL 6.4 (native Linux for System z)

Included is a recommended multipath.conf for Linux 6.x and 5.x hosts. Copy and paste the following text into a standard text file to use as the multipath.conf file.

Note: Be sure to uncomment the lines that pertain to the device and Linux host version for your system. In addition, be sure to modify the WWID line appropriately for your system.

```
defaults {
    udev_dir           /dev
    polling_interval   5
    checker_timeout    10
}
blacklist {
    wwid               "*"
}
blacklist_exceptions {
    wwid               "20020c24*"
}
multipaths {
    multipath {
        wwid           20020c24000093676
        alias          flashfs1
    }
}
devices {
    device {
        vendor          "IBM"
        product         "FlashSystem"
        path_selector   "round-robin 0"
        path_grouping_policy multibus
        path_checker    tur
#         rr_min_io_rq    1         # Linux 6.x, FlashSystem, RS 710/810
#         rr_min_io_rq    4         # Linux 6.x, FlashSystem, RS 720/820
#         rr_min_io       1         # Linux 5.x, FlashSystem, RS 710/810
#         rr_min_io       4         # Linux 5.x, FlashSystem, RS 720/820
        rr_weight       uniform
        no_path_retry   fail
        failback        immediate
        dev_loss_tmo    300
        fast_io_fail_tmo 25
    }
}
```

Verifying alignment using statistics tools

You can verify alignment by using the statistics tools on the storage system.

The storage system statistics can be monitored through the **Statistics** node in the system tree of the web management interface. You can view the statistics in real time or record them for later analysis.

The alignment-specific counters are logged at the port level. If Active/Active multipathing is employed, it is safe to review only one path. Be sure that the path in review presents the logical unit numbers (LUNs) to the application that is being monitored.

The counters that are displayed in the statistics that are related to this analysis are as follows:

- For storage system firmware 5.3.2, or later, refer to the *RMW Count* counter
- For firmware earlier than 5.3.2, the following counters apply:
 - 2k SCSI commands
 - 4k SCSI commands
 - 512 aligned SCSI commands
 - 2k aligned SCSI commands
 - 4k aligned SCSI commands
 - 1k aligned SCSI commands

Read-modify-writes (RMW) are the result of unaligned write operations, so for systems that support the RMW count statistic, this value is the only one needed to verify alignment.

2k SCSI commands and *4k SCSI commands* are counters for the size of the requests. The other four counters reference the starting address of the request. For an application to be considered aligned, the size of the request must be a multiple of 4 KB and the starting address of the request must be divisible by 4 KB.

To help the analysis, change the settings for the graph by clicking the **Configure** button in the **System Statistics** toolbar. Be sure to include the Delta and Plot Delta columns, and set the **Update Interval** to 5 seconds.

As the application runs, the counters increase. The counters are running counters and the Current column increases only across samples. With the Delta column, the number of commands that are received across samples is presented.

The RMW Count counter, the 2k SCSI commands counter, and the 4k aligned SCSI commands counter all reflect unaligned write operations. Because the objective of the application layout is to achieve alignment for maximum efficiency, the delta RMW Count counter value should be minimal. Most of the I/O activity should be accounted for by the 4k aligned counters.

Some unaligned I/O operations might be unavoidable, so do not expect 100% aligned accesses.

These analyzed counters can also be extended to include other counters at multiples of 4 KB.

If the application presents a high delta RMW Count counter or most of the I/O activity is on the 4 KB aligned counters, investigate the partition alignment for the operating system.

Notices

This information was developed for products and services offered in the U.S.A.

The manufacturer may not offer the products, services, or features discussed in this document in other countries. Consult the manufacturer's representative for information on the products and services currently available in your area. Any reference to the manufacturer's product, program, or service is not intended to state or imply that only that product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any intellectual property right of the manufacturer may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any product, program, or service.

The manufacturer may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to the manufacturer.

For license inquiries regarding double-byte character set (DBCS) information, contact the Intellectual Property Department in your country or send inquiries, in writing, to the manufacturer.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: THIS PUBLICATION IS PROVIDED "AS IS " WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. The manufacturer may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to websites not owned by the manufacturer are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this product and use of those websites is at your own risk.

The manufacturer may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact the manufacturer.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, IBM License Agreement for Machine Code, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the

same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning products not produced by this manufacturer was obtained from the suppliers of those products, their published announcements or other publicly available sources. This manufacturer has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to products not produced by this manufacturer. Questions on the capabilities of products not produced by this manufacturer should be addressed to the suppliers of those products.

All statements regarding the manufacturer's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

The manufacturer's prices shown are the manufacturer's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to the manufacturer, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. The manufacturer, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. The manufacturer shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

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

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

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal Use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative works of these publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.



Printed in USA