

The Lenovo logo is displayed in white text on a black rectangular background.

Understanding the Performance Benefits of MegaRAID FastPath with Lenovo Servers

Introduces the SSD Performance Accelerator for Lenovo servers

Describes the performance gains from using MegaRAID FastPath

Explains situations where FastPath should not be used

Shows lab measurements detailing the advantages of MegaRAID FastPath

Ray Engler
Beverly Clapp
Zach Purser



Abstract

MegaRAID FastPath is a feature of Lenovo® high-end RAID adapters. It enables the adapter to perform significantly better with solid-state drives (SSDs) in high IOPS situations. When FastPath is enabled in combination with arrays defined on fast storage, the maximum IOPS each RAID adapter can handle is significantly increased. FastPath has no effect on sequential workloads nor does it increase the performance of slow storage.

This paper discusses the FastPath feature, what it's designed to do, how to properly enable it, and what performance gains you should expect to get when using it.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

<http://lenovopress.com>

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

MegaRAID FastPath	3
Enabling FastPath	3
Using FastPath.....	6
Performance measurements	6
Summary	10
Change history.....	10
Authors.....	10
Notices	11
Trademarks	12

MegaRAID FastPath

MegaRAID FastPath is an optimization in the MegaRAID software stack on RAID controllers such as the Lenovo ServeRAID™ M5100 and M5200 series adapters on Lenovo System x® servers, as well as the ThinkServer® RAID 720i family of adapters on ThinkServer systems. FastPath is designed to be used with SSD arrays in high IOP situations such as OLTP database workloads.

The adapters listed in Table 1 support MegaRAID FastPath, either standard or as an option.

Table 1 *Lenovo RAID controllers that support MegaRAID FastPath*

Part number	Description
81Y4481	ServeRAID M5110 SAS/SATA Controller
81Y4478	ServeRAID M5120 SAS/SATA Controller
46C9110	ServeRAID M5210 SAS/SATA Controller
00AE938	ServeRAID M5225-2GB SAS/SATA Controller
4XC0G88831 4XC0G88849	ThinkServer RAID 720i PCIe Adapter
4XC0G88838	ThinkServer RAID 720i AnyRAID™ Adapter
4XC0G88839	ThinkServer RAID 720ix AnyRAID Adapter with Expander

With FastPath enabled, and when certain conditions are met, the MegaRAID software stack running on the RAID controller can bypass portions of the software stack, resulting in a highly efficient, streamlined code path that the RAID controller can execute very quickly. This means that the RAID controller can handle more trips through its code in a given amount of time, which equates to higher possible IOPS.

FastPath is not designed to speed up slow storage. It will not have an effect on storage performance unless the RAID controller is the bottleneck. If the controller itself is limiting storage performance, then enabling FastPath will allow the controller to run faster, raising performance. FastPath is not meant for use with HDD storage nor is it meant for sequential workloads.

Enabling FastPath

The FastPath feature is enabled in a variety of ways, depending on the adapter:

- ▶ With the ThinkServer RAID 720i adapters, FastPath is included with the adapter by default.
- ▶ With the ServeRAID M5100 and M5200 series controllers, the key is managed completely in software using the Lenovo Features on Demand (FoD) license upgrade process. The license for FastPath on these adapters can be purchased using one of the part numbers listed in the following table.

Table 2 *FastPath part numbers for ServeRAID M5100/M5200 controllers*

Part number	Description
90Y4273	M5100 Series Performance Accelerator

Part number	Description
47C8710	M5200 Series Performance Accelerator

The ServeRAID M5100/M5200 Series Performance Accelerator Key software license is installed in the server's Integrated Management Module (IMM). This is done by using the IMM's web interface, specifically the **IMM Management** → **Activation Key Management** screen, as shown in Figure 1.

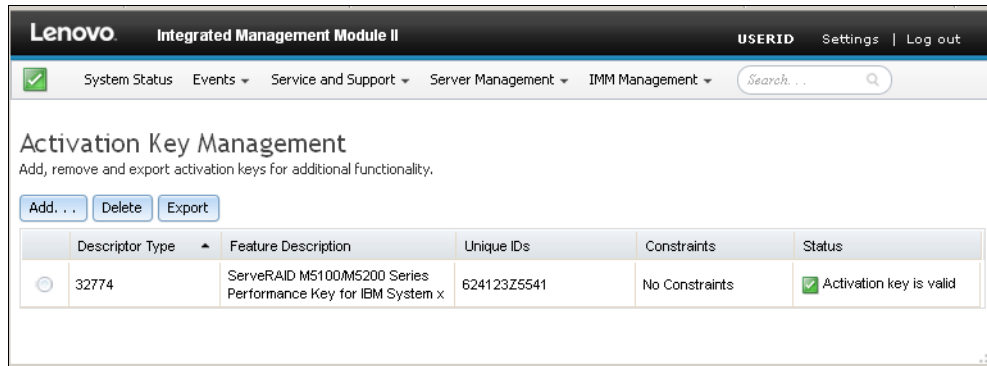


Figure 1 Performance Accelerator Key (MegaRAID FastPath) installed and active

Once the Performance Key is properly installed, the server must be rebooted for the key to take effect. FastPath will then be enabled on all of the ServeRAID M5100/M5200 RAID controllers installed in the system- one key enables FastPath on all of the installed M5100/M5200 adapters. The key is unique to each server and therefore cannot be moved from server to server. If more ServeRAID adapters are added to the server at a later time, FastPath will be enabled for all of them when the server is booted. If ServeRAID adapters are moved to another server, FastPath will only be enabled on them if that other server also has a Performance Accelerator Key installed in its IMM.

Once the server is booted, FastPath enablement can be verified by using storcli or megacli on the command line or by using MegaRAID Storage Manager (MSM). For example, Figure 2 shows how to use storcli to check controller 1:

```
C:\>storcli64 /c1 show aso
```

```
Controller = 1  
Status = Success  
Description = None
```

```
Premium Feature Key :
```

```
=====
```

```
-----  
Adv S/W Opt          Time Remaining  Mode  
-----
```

```
MegaRAID FastPath  Unlimited      -
```

```
MegaRAID RAID5     Unlimited      -  
-----
```

```
Re-host Information :
```

```
=====
```

```
Needs Re-hosting = NO
```

Figure 2 Using storcli to confirm FastPath is enabled

The presence of the Advanced Software Option “MegaRAID FastPath” (highlighted in red in the previous figure) verifies that FastPath is enabled for that specific ServerRAID M5100/M5200 controller. To verify it for all installed ServerRAID adapters, use this storcli command:

```
storcli64 /cALL show aso
```

The equivalent megacli command to check controller 1 is shown in Figure 3.

```
C:\>megacli64 elf controllerfeatures a1
```

```
Activated Advanced Software Options  
-----
```

```
Advanced Software Option      : MegaRAID FastPath
```

```
Time Remaining      : Unlimited
```

```
Advanced Software Option      : MegaRAID RAID5
```

```
Time Remaining      : Unlimited
```

```
Re-host Information  
-----
```

```
Needs Re-hosting : No
```

```
Exit Code: 0x00
```

Figure 3 Using megacli to confirm FastPath is enabled

Figure 4 shows the location of the FastPath status on the MegaRAID Storage Manager (MSM) dashboard screen.

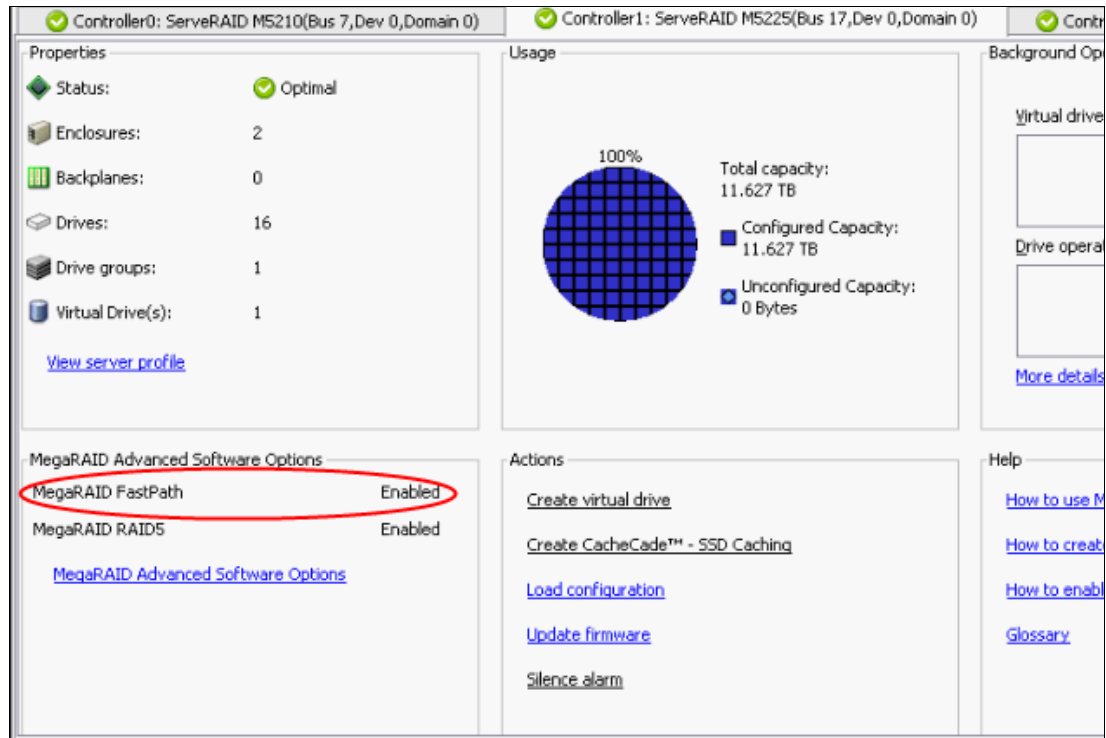


Figure 4 Location of FastPath status on the MSM dashboard

Using FastPath

To benefit from FastPath, an array must be defined using these specific parameters:

- ▶ Write Policy: Write Through
- ▶ IO Policy: Direct IO
- ▶ Read Policy: No Read Ahead

These parameters can be set as desired for FastPath arrays:

- ▶ Stripe Size: 64KB is the optimal size for FastPath for most workloads, though some workloads may perform better with other stripe sizes.
- ▶ RAID level: FastPath works with all RAID levels.
- ▶ Disk Cache Policy: This controls the write cache policy for the drives in the arrays (as opposed to the write cache policy of the RAID adapter). Enabling disk write caching can put data at risk.

When an array meeting the above requirements is used on a RAID controller with FastPath enabled, FastPath will be fully operational for that array.

Performance measurements

The following configuration was used to measure the benefits of FastPath:

- ▶ Lenovo System x3950 X6 server with 8x Intel Xeon E7 processors

- ▶ Windows Server 2012 R2 installed on an internal 2-drive RAID-1 HDD array connected to an M5210 SAS/SATA controller
- ▶ 16x 800GB 12Gb Enterprise-class SAS SSDs installed in two Lenovo D1024 6Gb SAS/SATA JBOD enclosures (eight SSDs per enclosure) both attached directly to one ServeRAID M5225 12Gb SAS/SATA controller (one enclosure per M5225 port, for best performance).
- ▶ One array was created using all sixteen of the SSDs, using these FastPath-compatible settings:
 - RAID-5
 - 64k stripe
 - no read-ahead
 - write through
 - direct IO
 - read/write
 - “unchanged” disk write cache
 - full initialization

Once the array initialization was complete, a partition was created and formatted NTFS. The measurements were taken using this partition. FastPath was toggled on or off by adding or removing the M5100/M5200 Performance Key in the system’s IMM and then rebooting the system.

Figure 5 shows the 4K 100% random read measurements as the queue depth is varied.

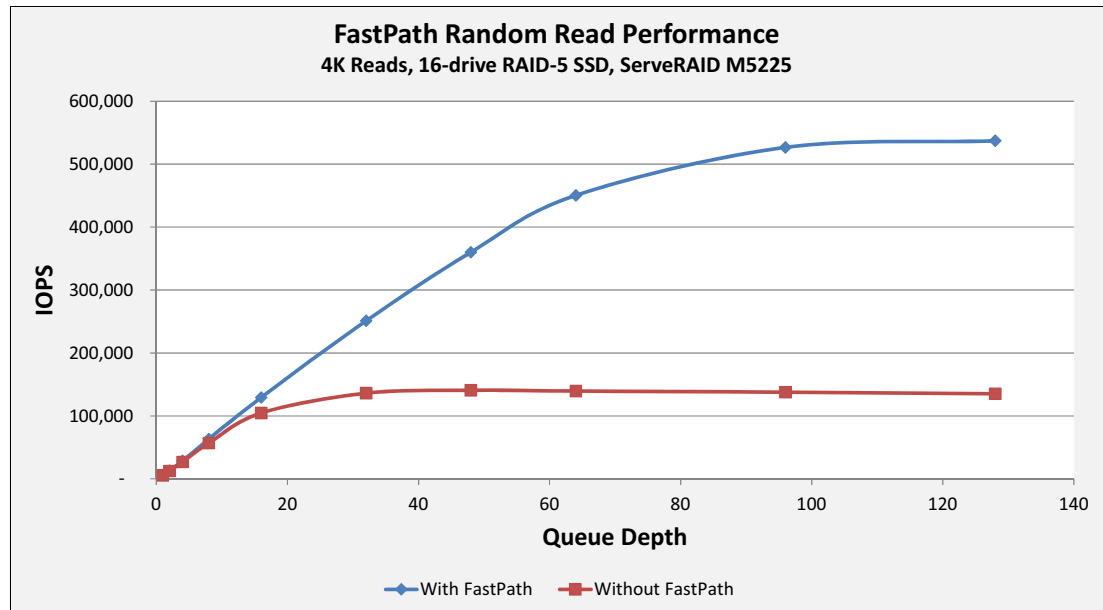


Figure 5 Comparison of FastPath enabled and disabled for 4K random read IOPS

On the left side of the graph, up to about 75,000 4K random read IOPS, performance is low because SSDs don’t perform very well at low queue depths. In that section of the graph, there really isn’t much of a difference there between FastPath being enabled or disabled because the controller can handle all the IOPS that the SSDs are putting out.

This negligible difference proves two points:

- ▶ FastPath’s impact is limited at low performance levels

- FastPath won't make slow storage faster- if the drives themselves are capable of only a few thousand IOPS, FastPath isn't going to help.

FastPath only helps when the drives are capable of delivering more IOPS than the RAID controller can handle. This is why FastPath isn't meant for HDD storage; it would take hundreds of HDDs to achieve 75,000 IOPS.

Once the queue depths increase, the SSDs reach more optimal performance and the effects of enabling FastPath are evident. With FastPath disabled, the controller tops out around 140,000 IOPS; with FastPath enabled the controller is able to handle up to nearly 550,000 4K random read IOPS, roughly 4 times the performance compared to FastPath being disabled! The effects of the streamlined FastPath code path in the controller are very dramatic.

Figure 6 shows the 4K 100% random write measurements as the queue depth is varied.

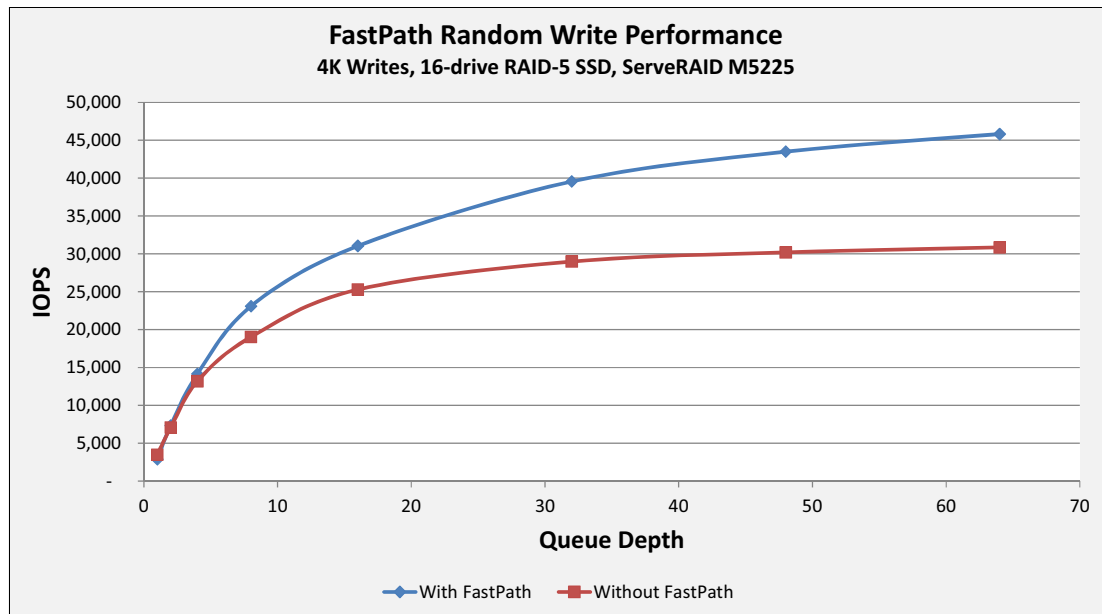


Figure 6 Comparison of FastPath enabled and disabled for 4K random write IOPS

As in the 100% random read case, FastPath is of little value when the IOPS are low, but it offers a significant performance advantage as the IOPS increase, starting at about 13,000 IOPS.

The absolute values of the IOPS in Figure 6 (100% writes) are much lower than in Figure 5 on page 7 (100% reads) due to the RAID 5 algorithm, which means for each IOPS shown in Figure 6 the controller is executing four IOPS with the SSDs (1 host-side RAID-5 write = 2 drive-side reads plus 2 drive-side writes). So 13,000 host-side RAID-5 write IOPS translates to 52,000 drive-side SSD IOPS.

Figure 7 on page 9 shows the measurements when the I/O pattern is 50% 4K random reads and 50% 4K random writes. Again, the benefits of FastPath are easily seen as the IOPS increase, starting at about 20,000 IOPS. The controller's streamlined FastPath code path allows it to operate much more efficiently.

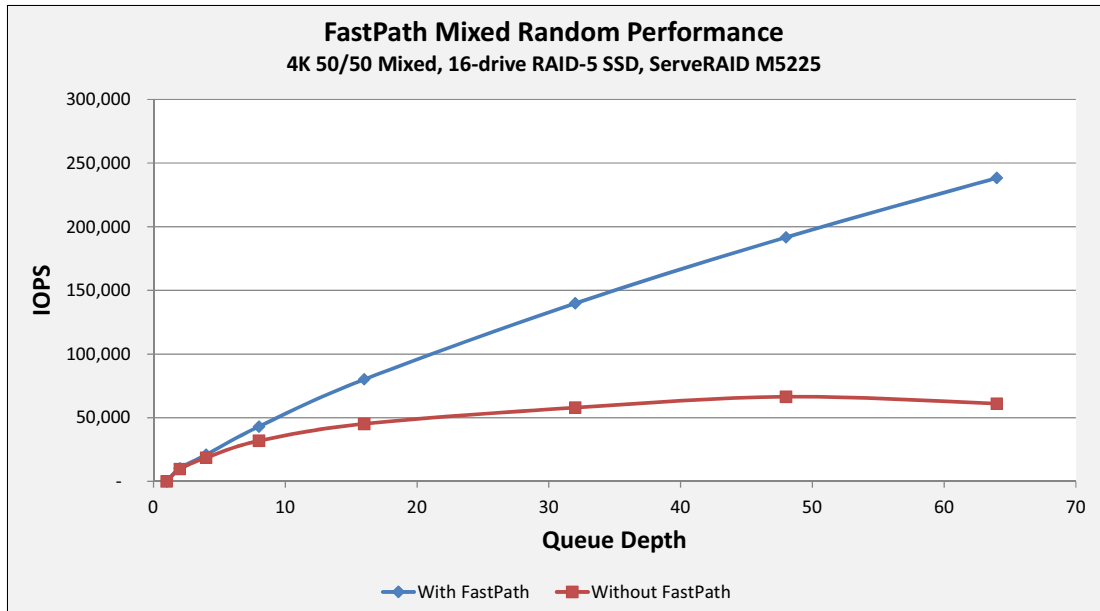


Figure 7 Comparison of FastPath enabled and disabled for 4K mixed 50/50 IOPS

Figure 8 shows the sequential read and write measurements. The difference between FastPath enabled and disabled is within the noise of the measurements; that is, there is effectively no difference. FastPath provides no benefit for sequential workloads. In that case, other tunings should be used, such as read-ahead, adapter write caching, and/or a larger stripe size.




Figure 8 Comparison of FastPath enabled and disabled for sequential reads and writes

Summary

MegaRAID FastPath is a feature of Lenovo's high performance RAID controllers that allows the controller to run much more efficiently in high IOPS situations, such as OLTP database workloads. Combined with a properly configured array defined on fast SSDs, using FastPath can result in significantly higher storage performance.

FastPath does not help storage performance in low IOPS situations, for sequential workloads, or if the underlying drives aren't capable of driving high IOPS (such as HDDs). FastPath will only provide a benefit when the drives are capable of handling more IOPS than the controller.

Change history

February 2017:

- ▶ Grammar and readability improvements

Authors

This paper was produced by the following team of specialists:

Ray Engler is an Advisory Hardware Engineer on the Lenovo Server Performance Team and is Lenovo's primary representative to the TPC (Transaction Processing Performance Council). He has worked with TPC benchmarks such as TPC-E and TPC-C for over 15 years, including the publication of dozens of #1 world-record benchmark results on configurations costing up to several million dollars. Ray joined Lenovo in October 2014 after having worked at IBM for 22 years. Prior to working on server performance, he worked on desktop performance, automation, and OS/2 development. Ray holds a Bachelor of Science degree in Computer and Systems Engineering from Rensselaer Polytechnic Institute in Troy, NY.

Beverly Clapp is an Advisory Hardware Performance Engineer with the Lenovo Server Performance Team in Raleigh, NC. Her focus is internal storage performance analysis on RAID adapters and HBAs using SAS/SATA HDDs and SSDs as well as flash-based solutions such as NVMe products. Beverly has a Bachelor's Degree in Computer Science from North Carolina A&T State University. She has 15 years of performance related experience with both workstations and servers.

Zach Purser is an Advisory Hardware Performance Engineer with the Lenovo Server Performance Team in Raleigh, NC. His focus is modeling RAID adapter and HBA performance with SAS/SATA HDDs and SSDs, and NVMe SSDs, and modeling server performance for general purpose workloads such as SPECcpu and similar benchmarks. Zach has a PhD in Processor Microarchitecture from North Carolina State University and has 12 years of server performance related experience.

Thanks to the following people for their contributions to this project:

- ▶ David Watts, Lenovo Press
- ▶ Mark T. Chapman, Lenovo

Notices

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

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

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "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 jurisdictions 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. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

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

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result 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.

This document was created or updated on February 3, 2017.

Send us your comments via the **Rate & Provide Feedback** form found at <http://lenovopress.com/1p0592>

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at <http://www.lenovo.com/legal/copytrade.html>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

AnyRAID™
Lenovo®

Lenovo(logo)®
ServeRAID™

System x®
ThinkServer®

The following terms are trademarks of other companies:

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

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

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