

Read Intensive Flash Drives

Shabbir Penkar



Storage

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get personalized notifications of new content
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK



Introduction

This IBM® Redpaper™ publication describes the configuration guidelines for using Read Intensive Flash Drives available from IBM with the IBM Spectrum Virtualize™ and IBM Storwize® family systems.

The Read Intensive (RI) solid state drives (SSDs) that are available on Storwize V7000 Gen2, Storwize V5000 Gen2, and IBM SAN Volume Controller 2145-DH8/24F are one Drive Write Per Day (DWPD) Read Intensive drives.

To deploy Read Intensive SSDs, your system must have Storwize code levels 7.6.1.4 or 7.7.0 installed.

RI SSDs are available to purchase as an add-on to an existing Storwize system, or they can be bought with a new system that contains RI SSDs. Due to the Read Intensive nature of these drives, they must be treated differently from regular or 10 DWPD SSD drives. We advise that the user follow the configuration guidelines that are detailed in the following document to be able to make the best use of RI SSDs.

Drive identification

RI SSDs on a system can be identified by searching for the `vendor_id` assigned to these drives. Depending on the Storwize model, the RI SSD will use one of the following `vendor_ids`:

- ▶ Storwize V7000 (Gen2): **IBM-D051**
- ▶ Storwize V5000 (Gen2): **IBM-E051**
- ▶ IBM SAN Volume Controller 2145-DH8/24F: **IBM-C051**

For example, on a Storwize V7000 Gen2 system, the RI SSD drives can be identified by using the `lsdrive` command, as shown in Example 1.

Example 1 The lsdrive command

```
IBM_Storwize:cluster_name:admin>lsdrive -gui |grep -i d051
15 onlinecandidate 50000396ec8bbb91 sas_hdd1.7TB512IBM-D051 PX04SRB19201EJ595
11S00VN228YXXXSVQARGWJ 10000 540D27onlineonlineno12Gbinactive 129
23 onlinecandidate 50000396ec8bbbe5 sas_hdd1.7TB512IBM-D051 PX04SRB19201EJ595
11S00VN228YXXXSVQARGZ6 10000 540D26onlineonlineno12Gbinactive 129
```

Using the Storwize graphical user interface (GUI), you can also view the RI SSD pools by navigating to **Pools** → **Internal Storage**, as shown in Figure 1.

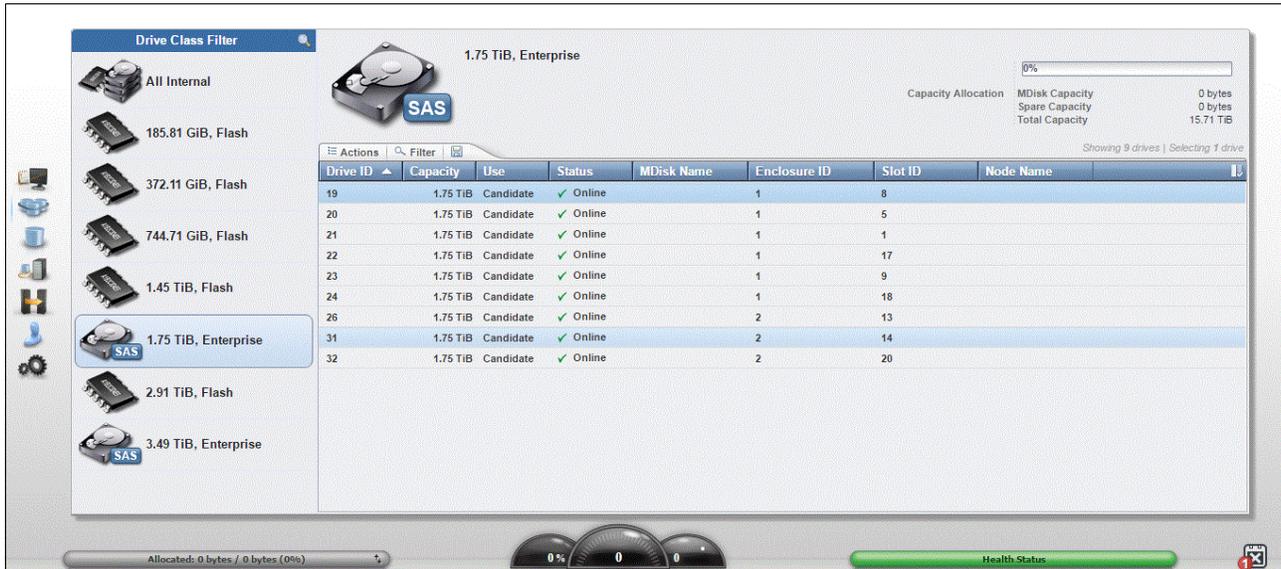


Figure 1 RI SSD pools

Note that RI SSDs will have the following attributes when listed by using `lsdrive <drive_id>` or `lsdrive -gui` or in the Properties view using the GUI, as shown in Example 2.

Example 2 RI SSD attributes

```
product_id: PX04SRB192
vendor_id: IBM-D051/IBM-E051/IBM-C051 (depending on Storwize product)
tech_type: sas_hdd
RPM: 10000
FRU_part_number: 01EJ601
```

For example, a 2 terabyte (TB) RI SSD on a Storwize V7000 Gen2 system will display the information shown in Example 3.

Example 3 The lsdrive <drive_id> command

```
IBM_Storwize:cluster_name:admin>lsdrive 99
id 99
status online
error_sequence_number
use candidate
UID 50000396ec8bbbbd
tech_type sas_hdd
capacity 1.7TB
block_size 512
vendor_id IBM-D051
product_id PX04SRB192
FRU_part_number 01EJ595
FRU_identity 11S00VN228YXXSVQARGXW
RPM 10000
firmware_level 540D
FPGA_level
mdisk_id
```

```
mdisk_name
member_id
enclosure_id 6
slot_id 24
node_id
node_name
quorum_id
port_1_status online
port_2_status online
interface_speed 12Gb
protection_enabled yes
auto_manage inactive
drive_class_id 129
```

Figure 2 shows the display that uses the GUI Properties view for the drive, and the Vendor ID indicates an RI SSD.

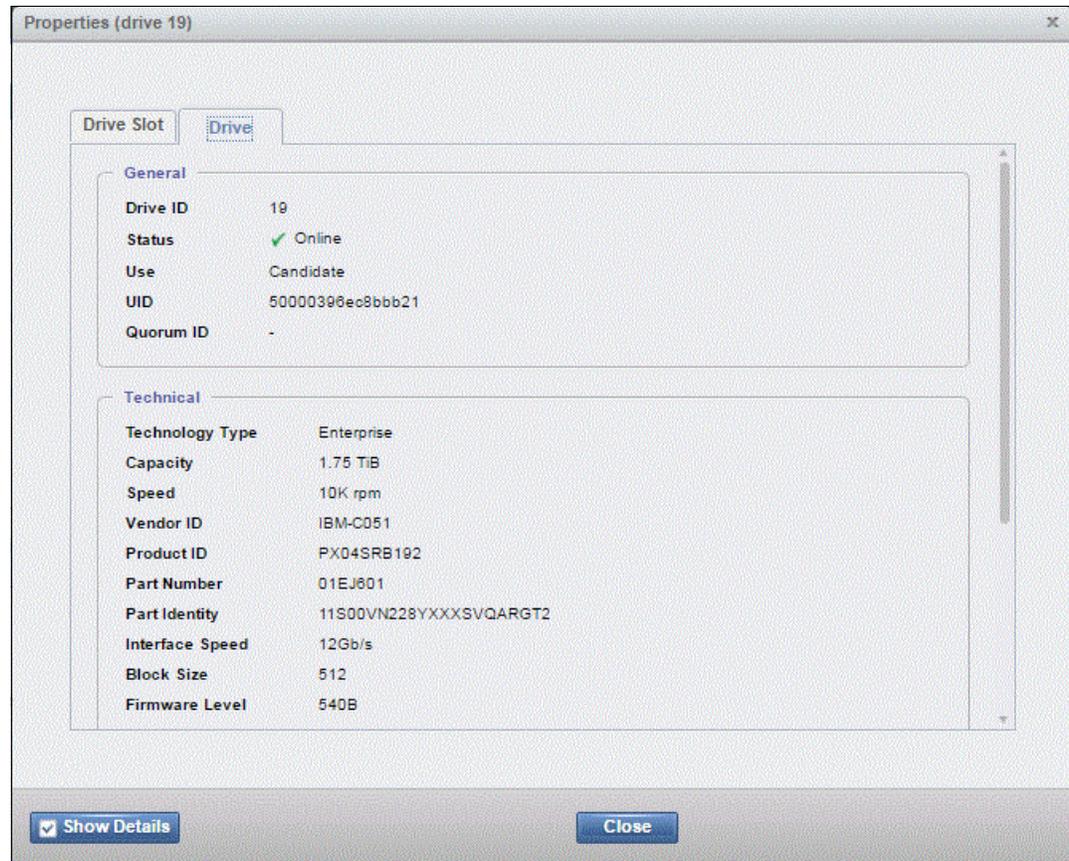


Figure 2 Vendor ID indicates RI SSD

RAID considerations

In this current (at the time of writing) release of Storwize code (version 7.7.0/7.6.1.4), RI SSDs are assigned a `tech_type` of `sas_hdd` at 10,000 revolutions per minute (RPM). Due to this fact and the limitations on the write capability of these SSDs, the following leading practices apply when you create Redundant Array of Independent Disks (RAID) configurations:

- ▶ It is required that Read Intensive SSDs RAID and managed disk (MDisk) pools are created manually.
- ▶ If you are configuring a new system, create RI SSD RAIDs and pools before you configure other drives (other SSDs, Enterprise and Nearline serial-attached SCSI (SAS) drives).
- ▶ If you are adding RI SSDs into an existing system, create RAID and storage pools manually. Do not mix RI SSD and any other drive type in the same pool.
- ▶ Given that RI SSDs are assigned a `tech_type` of `sas_hdd` and 10,000 RPM, it is possible that automatic configuration tools might mix these SSDs with other 10,000 RPM hard disk drives (HDDs). This can affect RI SSD array performance. It is strongly advised that automatic storage configuration tools are *not* used for these SSDs.

It is imperative that enough spare RI SSDs are made available for RI SSD arrays. If a drive in an RI SSD array fails and a matching RI SSD is not available as a spare, the system picks up a HDD of matching or higher capacity and incorporates it into the RI SSD array. This impairs the performance of the entire RI SSD array. It is essential that sufficient spare RI SSDs are available to maintain exclusivity of RI SSD arrays.

Figure 3 and Figure 4 illustrate leaving sufficient null members of spares for RI arrays.

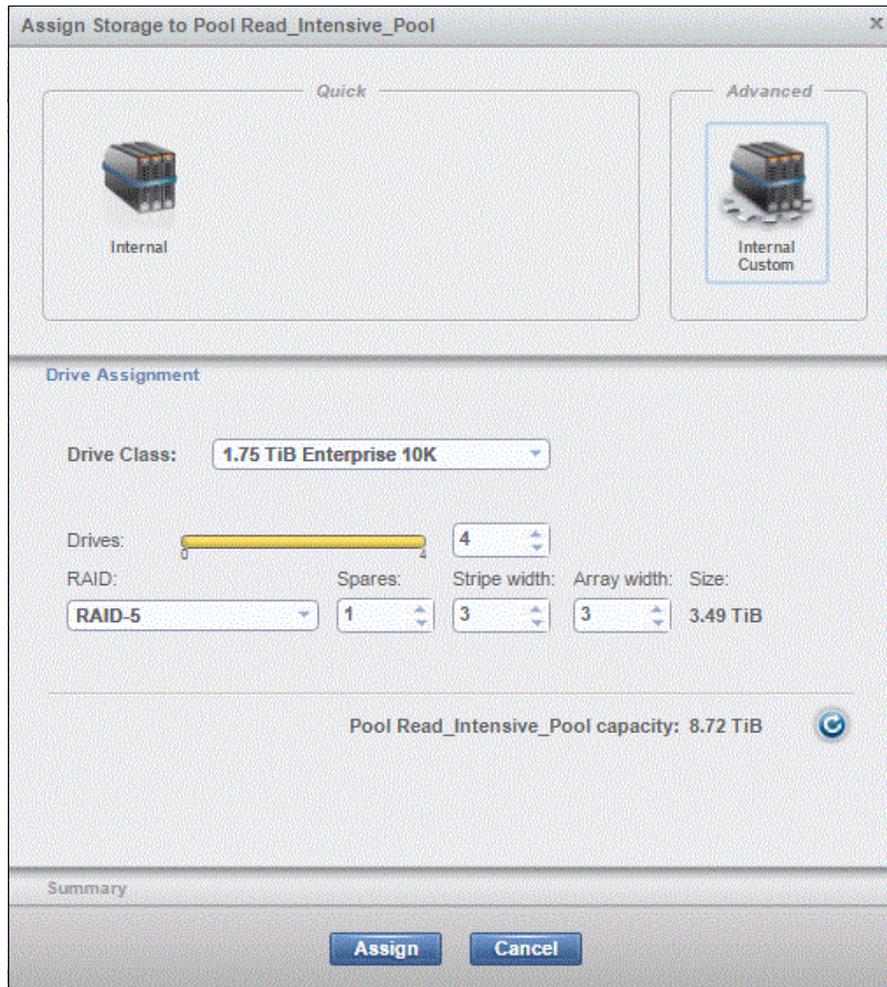


Figure 3 Spares



Figure 4 Spares (2)

It is necessary to ensure that RAID arrays formed of HDDs are assigned enough spare HDDs. In the absence of this configuration, an RI SSD could get incorporated into an HDD RAID array, severely under-utilizing the RI SSD.

RI SSD MDisks should not be mixed with other HDD MDisks in the same pool. If incorporated into a hybrid pool, IBM Easy Tier® management could subject the RI SSDs to unsustainable workloads.

Read Intensive SSD considerations

RI drives have sufficient endurance to permit an average of one DWPD, and clients should choose applications for RI SSDs that have a greater read versus write workload.

When an RI SSD is used up to 95% of its endurance, an alert is logged in the system event log. If the SSD is a redundant drive in a RAID array, it will be taken offline. If it is non-redundant (taking this SSD offline will result in an array going offline), the SSD continues to be available. It is advised that you arrange a replacement as soon as possible after receiving this alert.

The eventlog entry has an event ID of 010073 and a description of Drive reporting PFA Errors, as shown in Example 4 and Figure 5 on page 7.

Example 4 PFA errors (lines removed for brevity)

```
IBM_Storwize:cluster_name:admin>|seventlog 161
sequence_number 161
notification_type error
event_id 010073
event_id_text Drive reporting PFA Errors
error_code 1680
error_code_text Drive fault type 1
```

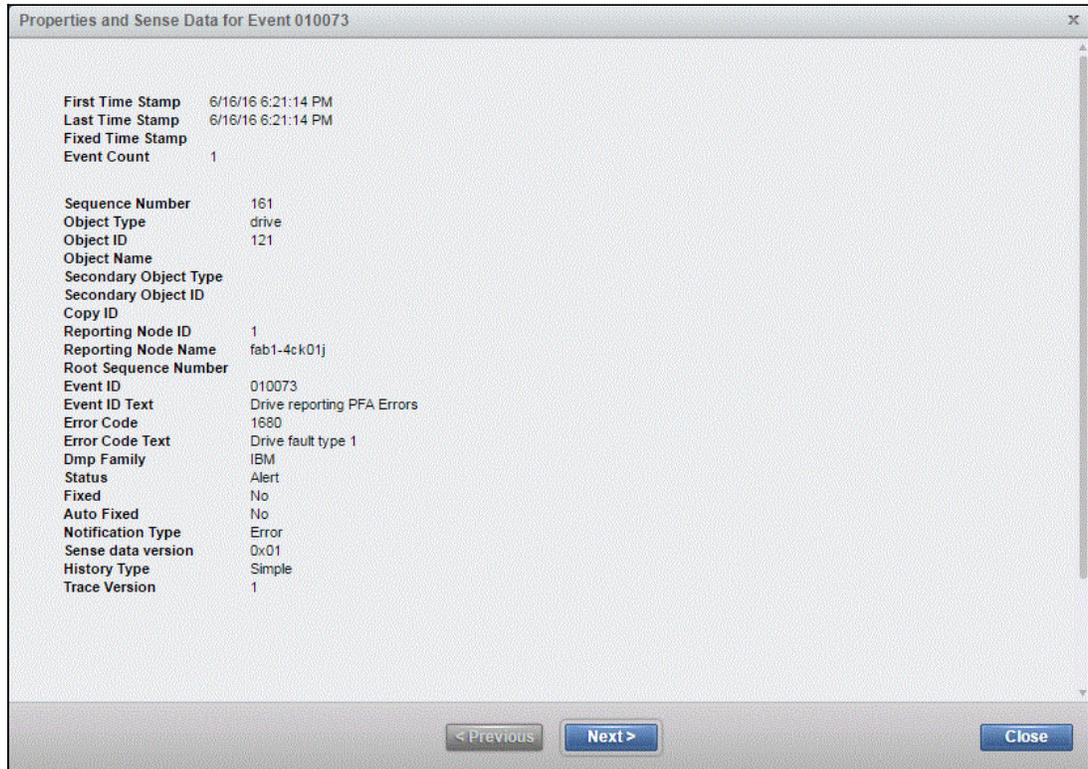


Figure 5 PFA errors

System resource use impact on endurance

Keep in mind that there will also be some system resources expended using RI SSD. For example, array initialization will write zeros to all drives in the RAID array. If the drives are used as quorum disks, system data is written to a portion of the disks in addition to user data. The impact on endurance from these operations is considered minimal.

When volumes (VDisks) are created, disk space occupied by the VDisks are initialized with zeros. Again, although impact to endurance is minimal, it is advised that you create all volumes by using the `mkvdisk -nofmtdisk` command, as shown in Example 5.

Example 5 The mkvdisk command

```
IBM_Storwise:cluster_name:admin>mkvdisk -nofmtdisk <other options>
```

Authors

This paper was produced at IBM Hursley UK labs.

Shabbir Penkar has worked in the storage industry for the last 15 years. He works at IBM Hursley Labs, and is responsible for certifying HDDs and SSDs on various Storwize products.

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

Sukhi Sohal

Suri Poliseti

John Fairhurst

IBM Hursley

Now you can become a published author, too

Here's an opportunity to spotlight your skills, grow your career, and become a published author, all at the same time. Join an International Technical Support Organization (ITSO) residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies.

Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run 2 - 6 weeks in length, and you can participate either in person or as a remote resident working from your home base.

Learn more about the residency program, browse the residency index, and apply online:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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

IBM 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:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION 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. IBM 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 non-IBM websites 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 IBM product and use of those websites is at your own risk.

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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 IBM, 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. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Easy Tier®
IBM®
IBM Spectrum™

Redbooks®
Redpaper™
Redbooks (logo) ®

Storwize®

The following terms are trademarks of other companies:

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



REDP-5380-00

ISBN 073845544X

Printed in U.S.A.

Get connected

