



Chrisanthy Carlane
Alex Osuna

IBM System Storage N series Implementation of RAID Double Parity for Data Protection

Overview

This IBM® Redpaper discusses the implementation of RAID double parity in IBM System Storage™ N series.

The following topics are covered:

- ▶ The need for RAID-DP
- ▶ How RAID-DP works
- ▶ Protection levels with RAID-DP

Abstract

To date, RAID technology, with its single parity, offered protection from a single failed disk with the caveat that no other disk fail or, increasingly more likely as disk media grow larger, an uncorrectable bit error would not occur during a read operation, while reconstruction of data lost on the failed disk was still in progress.

If either secondary event occurs during reconstruction, then some or all data contained in the RAID array could be lost.

IBM Storage System N series introduces double-parity RAID, named RAID-DP. This document provides an overview of RAID-DP and how it dramatically increases data fault tolerance from various disk failure scenarios. Other key areas covered include how much RAID-DP costs, special hardware requirements, and converting from existing RAID4-based volumes to RAID-DP. This document presents a double-disk failure recovery scenario to show how RAID-DP both allows the volume to continue serving data and recreates data lost on the two failed disks.

The need for RAID-DP

As mentioned earlier, traditional single-parity RAID offers adequate protection against a single event, which could be either a complete disk failure or a bit error during a read. In either event, data is recreated using both parity and data remaining on unaffected disks in the array or volume. If the event is a read error, then recreating data happens almost instantaneously, and the array or volume remains in an online mode. However, if a disk fails, then all data lost on it has to be recreated, and the array or volume will remain in a vulnerable degraded mode until data has been reconstructed onto a spare disk. It is in degraded mode that traditional single-parity RAID shows its protection capabilities have not kept up with modern disk architectures. RAID Double Parity significantly increases the fault tolerance from failed disk drives over traditional RAID.

The effect of modern larger disk sizes on RAID

Modern disk architectures have continued to evolve, as have other computer-related technologies. Disk drives are orders of magnitude larger than they were when RAID was first introduced. As disk drives have gotten larger, their reliability has not improved, and, more importantly, the bit error likelihood per drive has increased proportionally with the larger media. These three factors—larger disks, unimproved reliability, and increased bit errors with larger media—all have serious consequences for the ability of single-parity RAID to protect data.

Given that disks are as likely to fail now as when RAID technology was first introduced to protect data from such an event, RAID is still as vital now as it was then. When one disk fails, RAID simply recreates data from both parity and the remaining disks in the array or volume onto a spare disk. However, since RAID was introduced, the significant increases in disk size have resulted in much longer reconstruct times for data lost on the failed disk. Simply put, it takes much

longer to recreate data lost when a 274 GB disk fails than when a 36 GB disk fails. See Figure 1 for an illustration.

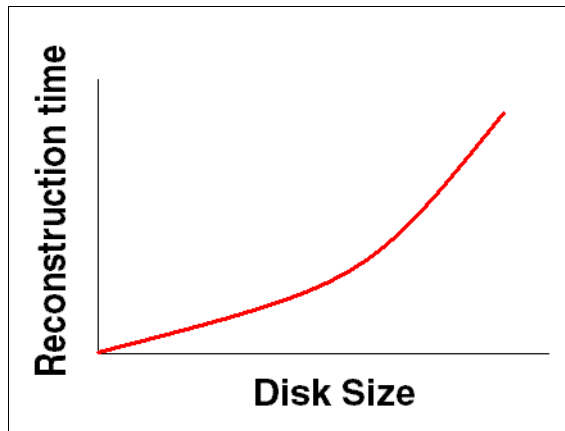


Figure 1 Reconstruction time versus disk size

Protection schemes with single-parity RAID using larger disks

The various options to extend the ability of single-parity RAID to protect data as disks continue to get larger are not attractive. The first is to continue to buy and implement storage using the smallest disk sizes possible so that reconstruction after a failed disk completes more quickly. However, this approach isn't practical from any point of view. Capacity density is critical in space constrained data centers, and smaller disks result in less capacity per square foot. Moreover, storage vendors are forced to offer products based on what disk manufacturers are supplying, and smaller disks aren't readily available, if at all. The second way to protect data on larger disks with single-parity RAID is slightly more practical but, with the introduction of RAID-DP, a less attractive approach for various reasons. Namely, by keeping the size of RAID arrays small, the time to reconstruct is reduced as seen in Figure 1. Continuing the analogy about a larger disk taking longer to reconstruct than a smaller one, a RAID array built with more disks takes longer to reconstruct data from one failed disk than one built with fewer disks. However, smaller RAID arrays have two costs that cannot be overcome. The first cost is that additional disks will be lost to parity impacting usable capacity and total cost of ownership (TCO). The second cost is that performance is generally slower with small RAID arrays, impacting business and users.

The most reliable protection offered by single-parity RAID is RAID1 or mirroring. See Figure 2 for an illustration. In RAID1, the mirroring process replicates an exact copy of all data on an array to a second array. While RAID1 mirroring affords maximum fault tolerance from disk failure, the cost of the implementation

is severe, since it takes twice the disk capacity to store the same amount of data. Earlier it was mentioned that using smaller RAID arrays to improve fault tolerance increases the total cost of ownership of storage due to less usable capacity per dollar spent. Continuing this approach, RAID1 mirroring, with its unpleasant requirement for double the amount of capacity, is the most expensive type of storage solution with the highest total cost of ownership.

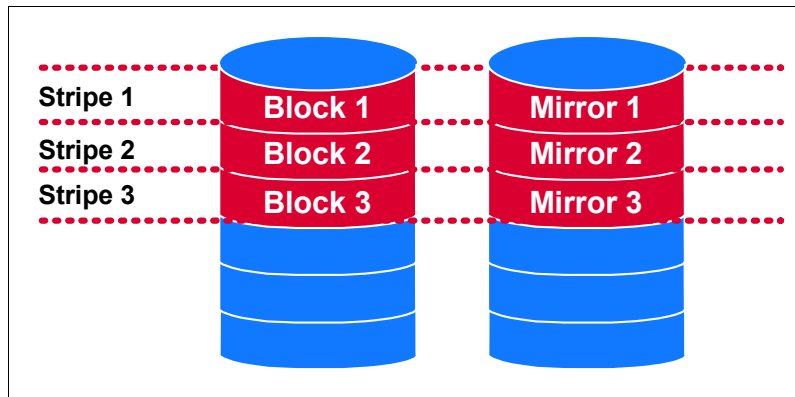


Figure 2 RAID1 protection

RAID-DP data protection

In short, given the current landscape, with larger disk drives affecting data protection, customers and analysts demand a better story about affordably improving RAID reliability from storage vendors. To meet this demand, IBM System Storage N series uses a new type of RAID protection named RAID-DP. RAID-DP stands for RAID Double Parity, and it significantly increases the fault tolerance from failed disk drives over traditional RAID. When all relevant numbers are plugged into the standard mean time to data loss (MTTDL) formula for RAID-DP versus single-parity RAID, RAID-DP is more reliable on the same underlying disk drives. With this reliability, RAID-DP approaches RAID1 mirroring for fault tolerance, but at RAID4 pricing. RAID-DP offers businesses the most compelling total cost of ownership storage option without putting their data at an increased risk.

How RAID-DP works

Traditional levels of existing RAID technology offer data protection through various approaches. The RAID-DP used by IBM System Storage N series, a modified RAID4, stores data in horizontal rows, calculates parity for data in the row, then stores the parity in a separate row parity disk. However, a constant

across the different RAID levels, including the modified RAID4, was that they used a single-parity scheme, which in turn limits their ability to protect past a single disk failure.

RAID-DP with double parity

It is well known that parity generally improves fault tolerance and that single-parity RAID improves data protection. Given that traditional single-parity RAID has established a very good track record to date, the concept of double-parity RAID should certainly sound like a better protection scheme. This validates the earlier example using the MTTFDL formula. But what exactly is RAID-DP with its double parity?

At the most basic layer, RAID-DP adds a second parity disk to each RAID group in an aggregate or traditional volume. A RAID group is an underlying construct that aggregates and traditional volumes are built upon. Each traditional RAID4 group has some number of data disks and one parity disk, with aggregates and volumes containing one or more RAID4 groups. Whereas the parity disk in a RAID4 volume stores row parity across the disks in a RAID4 group, the additional RAID-DP parity disk stores diagonal parity across the disks in a RAID-DP group. With these two parity stripes in RAID-DP, one the traditional horizontal, and the other diagonal, data protection is obtained even in the event of two disk failure events occurring in the same RAID group. See Figure 3 for a representation.

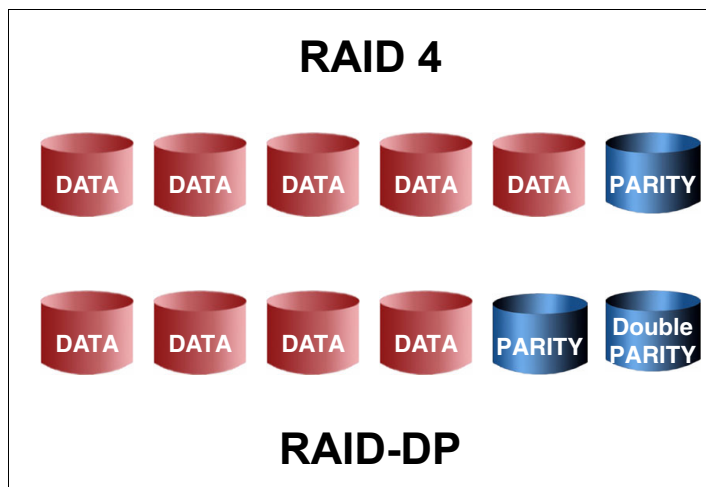


Figure 3 RAID4 and RAID-DP comparison

So, with RAID-DP, the traditional RAID4 horizontal parity structure is still employed and becomes subset of the RAID-DP construct. In other words, how RAID4 works on IBM System Storage N series has not been modified with

RAID-DP. The same process, in which data is written out in horizontal rows with parity calculated for each row, still holds in RAID-DP and is considered the row component double parity. In fact, if a single disk fails or a read error from a bad block or bit error occurs, then the row parity approach of RAID4 is the sole vehicle used to recreate the data without ever engaging RAID-DP. In this case, the diagonal parity component of RAID-DP is simply a protective envelope around the row parity component.

RAID4 horizontal row parity

Figure 4 illustrates the horizontal row parity approach used in the traditional RAID4 solution and is the first step in establishing an understanding of RAID-DP and double parity.




Figure 4 RAID4 horizontal row parity

The diagram represents a traditional RAID4 group using row parity that consists of four data disks (the first four columns, labeled “D”) and the single row parity disk (the last column, labeled “P”). The rows in Figure 4 represent the standard 4 KB blocks used by the traditional RAID4 implementation. The second row in the above diagram has been populated with some sample data in each 4 KB block and parity calculated for data in the row then stored in the corresponding block on the parity disk. In this case, the way parity was calculated was to add the values in each of the horizontal blocks, then store the sum as the parity value ($3 + 1 + 2 + 3 = 9$). In practice, parity is calculated by an exclusive OR (XOR) process, but addition is fairly similar and works as well for the purposes of this example. If the need arose to reconstruct data from a single failure, the process used to generate parity would simply be reversed. For example, if the first disk were to fail, when RAID4 recreated the data value 3 in the first column above, it would subtract the values on remaining disks from what is stored in parity ($9 - 3 - 2 - 1 = 3$). This example of reconstruction with single-parity RAID should

further assist with the conceptual understanding of why data is protected up to but not beyond one disk failure event.

Adding RAID-DP double parity stripes

In Figure 5, one diagonal parity stripe is added, denoted by the blue-shaded blocks, and a second parity disk, denoted with a “DP” in the sixth column, to the existing RAID4 group from the previous section and shows the RAID-DP construct that is a super-set of the underlying RAID4 horizontal row parity solution.




Figure 5 A set of RAID double parity stripes

The diagonal parity stripe has been calculated using the addition approach for this example rather than the XOR used in practice as discussed earlier and stored on the second parity disk ($1 + 2 + 2 + 7 = 12$). One of the most important items to note at this time is that the diagonal parity stripe includes an element from row parity as part of its diagonal parity sum. RAID-DP treats all disks in the original RAID4 construct, including both data and row parity disks, as the same.

Figure 6 adds in the rest of the data for each block and creates corresponding row and diagonal parity stripes.




Figure 6 Blocks of RAID double parity stripes

One RAID-DP condition that is apparent from Figure 6 is that the diagonal stripes wrap at the edges of the row parity construct.

Two important conditions for RAID-DP's ability to recover from double disk failures:

- ▶ Each diagonal parity stripe misses one and only one disk, but each diagonal misses a different disk.
- ▶ There is one diagonal stripe that doesn't get parity generated on it or get stored on the second diagonal parity disk. In this example the omitted diagonal stripe is the white noncolored blocks.

In the reconstruction example that follows, it will be apparent that omitting the one diagonal stripe doesn't affect RAID-DP's ability to recover all data in a double disk failure.

It is important to note that the same RAID-DP diagonal parity conditions covered in this example hold in real storage deployments that involve dozens of disks in a RAID group and millions of rows of data written horizontally across the RAID4 group. And, while it is easier to illustrate RAID-DP with the smaller example above, recovery of larger size RAID groups works exactly the same regardless of the number of disks in the RAID group.

Proving that RAID-DP really does recover all data in the event of a double disk failure can be done in two manners. One is using mathematical theorems and proofs, and the other is to simply go through a double disk failure and subsequent recovery process. This document will use the latter approach to prove the concept of RAID-DP double-parity protection. For more extensive coverage on the mathematical theorems and proofs of what RAID-DP is built

upon, please review the Row-Diagonal Parity for Double Disk Failure Correction document available at the USENIX Organization Web site:

<http://www.usenix.org>

RAID-DP reconstruction

See Figure 7 as the starting point for the double disk failure, and assume that the RAID group is functioning normally when a double disk failure occurs. This is denoted by all data in the first two columns are now missing in Figure 7.

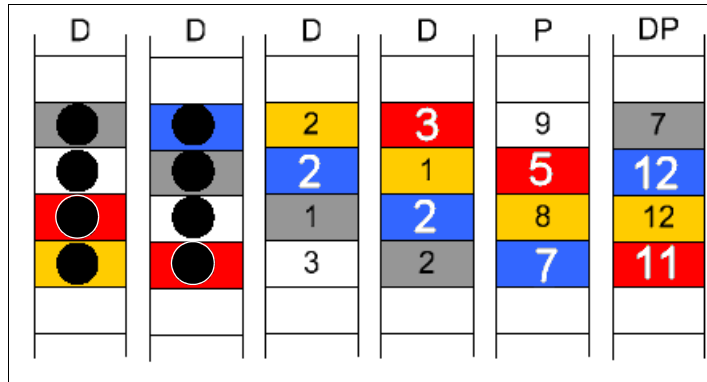


Figure 7 RAID double parity simulation of double disk failure

When engaged after a double disk failure, RAID-DP first begins looking for a chain to start reconstruction on. In this case, let's say the first diagonal parity stripe in the chain it finds is represented by the blue diagonal stripe. Remember when reconstructing data for a single disk failure under RAID4 that this is possible if and only if no more than one element is missing. With this in mind, traverse the blue diagonal stripe in the above diagram and notice that only one of the five blue blocks is missing. With four out of five elements available, RAID-DP has all of the information needed to reconstruct the data in the missing blue block. Figure 8 reflects this data having been recovered onto an available hot spare disk.




Figure 8 RAID double parity reconstruction simulation of blue diagonal block

The data has been recreated from the missing blue diagonal block using the same arithmetic discussed earlier ($12 - 7 - 2 - 2 = 1$). Now that the missing blue diagonal information has been recreated, the recovery process switches from using diagonal parity to using horizontal row parity. Specifically, in the top row after the blue diagonal has recreated the missing diagonal block, there is now enough information available to reconstruct the single missing horizontal block from row parity ($9 - 3 - 2 - 1 = 3$). See Figure 9.




Figure 9 RAID double parity reconstruction simulation of the first horizontal block

RAID-DP next continues in the same chain to determine if other diagonal stripes can be recreated. With the top left block having been recreated from row parity, RAID-DP can now recreate the missing diagonal block in the gray diagonal stripe, as follows. See Figure 10.




Figure 10 RAID double parity reconstruction simulation of gray diagonal block

Once again, after RAID-DP has recovered a missing diagonal block in a diagonal stripe, enough information exists for row parity to recreate the one missing horizontal block in the first column, as illustrated in Figure 11.




Figure 11 RAID double parity reconstruction simulation of the second horizontal block

As we noted earlier, the white diagonal stripe is not stored, and no additional diagonal blocks can be recreated on the existing chain. RAID-DP will start to look for a new chain to start recreating diagonal blocks on, and, for the purposes of this example, determines it can recreate missing data in the gold diagonal stripe, as Figure 12 shows.




Figure 12 RAID double parity reconstruction simulation of gray diagonal block

After RAID-DP has recreated a missing diagonal block, the process again switches to recreating a missing horizontal block from row parity. When the missing diagonal block in the gold diagonal has been recreated, enough information is available to recreate the missing horizontal block from row parity, as evident in Figure 13.




Figure 13 RAID double parity reconstruction simulation of the last horizontal block

After the missing block in the horizontal row has been recreated, reconstruction switches back to diagonal parity to recreate a missing diagonal block. RAID-DP can continue in the current chain on the red diagonal stripe, as shown in Figure 14.




Figure 14 RAID double parity reconstruction simulation of red diagonal block

Once again, after the recovery of a diagonal block the process switches back to row parity, as it has enough information to recreate data for the one horizontal block. The final diagram in the double disk failure scenario follows next, with all data having been recreated with RAID-DP. See Figure 15.




Figure 15 RAID double parity reconstruction simulation with all blocks recovered

Protection levels with RAID-DP

The previous recovery example goes a long way to give a pictorial description of RAID-DP in operation. But there are a few more areas about RAID-DP operations the example didn't make evident, but that need further discussion. If a double disk failure occurs, RAID-DP automatically raises the priority of the reconstruction process so the recovery completes faster. As a result, the time to reconstruct data from two failed disks is slightly less than the time to reconstruct data from a single disk failure. A second key feature of RAID-DP with double disk

failure is that it is highly likely one disk failed some time before the second and at least some information has already been recreated with traditional row parity. RAID-DP automatically adjusts for this occurrence by starting recovery where two elements are missing from the second disk failure.

A higher level of protection is available by using RAID-DP in conjunction with SyncMirror. In this configuration, the protection level is up to five concurrent disk failures, four concurrent disk failures followed by a bad block or bit error before reconstruction is completed.

Creating RAID-DP aggregates and traditional volumes

To create an aggregate or traditional volume with RAID-DP–based RAID groups, select the option in FilerView when provisioning storage or add the `-t raid_dp` switch to the traditional `aggr` or `vol` create command on the command line interface. The command line interface syntax would be `[vol | aggr] create name -t raid_dp X` (with X representing the number of disks the traditional volume or aggregate contains). If the type of RAID group is not specified, Data ONTAP will automatically use the default RAID group type. The default RAID group type used, either RAID-DP or RAID4, depends on the platform and disk being used. Please refer to the following table for the RAID group type defaults in Data ONTAP 7.1 and later.

Table 1 Data ONTAP 7.1 default RAID group type

	FC Disk	SATA Disk
Traditional Volumes	RAID4	RAID-DP
Aggregates with Flexible Volumes	RAID-DP	RAID-DP

The following output from the `vol status` command shows a four-disk RAID-DP RAID group for a traditional volume named `test`, with the second parity disk for diagonal parity denoted as *dparity*.

```

Volume test (online, raid_dp) (zoned checksums)
Plex /test/plex0 (online, normal, active)
RAID group /test/plex0/rg0 (normal)
RAID Disk Device HA SHELF BAY CHAN Used (MB/blks) Phys (MB/blks)
-----
dparity   v0.2  v1   0   1  FC:A 36/74752    42/87168
parity    v0.3  v0   0   3  FC:A 36/74752    42/87168
data      v0.6  v1   0   2  FC:A 36/74752    42/87168
data      v0.4  v1   0   4  FC:A 36/74752    42/87168

```

Figure 16 Sample of volume status command output

Converting existing aggregates and traditional volumes to RAID-DP

Existing aggregates and traditional volumes are easily converted to RAID-DP using the command `[aggr | vol] options name raidtype raid_dp`. See Figure 17 and Figure 18. When entered, the aggregate or traditional volume is instantly denoted as RAID-DP, but all diagonal parity stripes still need to be calculated and stored on the second parity disk. RAID-DP protection against double disk failure isn't available until all diagonal parity stripes have been calculated and stored on the diagonal parity disk. Calculating the diagonals as part of a conversion to RAID-DP takes time and impacts performance slightly on the storage controller. The amount of time and performance impact for conversions to RAID-DP depends on what the storage controller is and how busy the storage controller is during the conversion. Generally, conversions to RAID-DP should be planned for off-peak hours to minimize potential performance impact to business or users. For conversions from RAID4 to RAID-DP, certain conditions are required. Namely, conversions occur at the aggregate or traditional volume level, and there has to be an available disk for the second diagonal parity disk for each RAID4 group in either class of storage. The size of the disks used for diagonal parity need to be at least as big as the original RAID4 row parity disks.

```

node1> aggr status aggr1
Aggr State      Status          Options
aggr1 online    raid4, aggr

Volumes: voll
Plex /aggr1/plex0: online, normal, active
RAID group /aggr1/plex0/rg0: normal

node1> aggr options aggr1 raidtype raid_dp
Mon Mar 20 16:29:22 GMT [raid.config.raidsz.chng.notice]: aggregate
aggr1: raidsize is adjusted from 8 to 16 after changing raidtype
Mon Mar 20 16:29:39 GMT [raid.rg.recons.missing.notice]: RAID group
/aggr1/plex0/rg0 is missing 1 disk(s).
Mon Mar 20 16:29:39 GMT [raid.rg.recons.info.notice]: Spare disk v1.16 will
be used to reconstruct one missing disk in RAID group /aggr1/plex0/rg0.
Mon Mar 20 16:29:39 GMT [raid.rg.recons.start.notice]: /aggr1/plex0/rg0:
starting reconstruction, using disk v1.16
Mon Mar 20 16:29:53 GMT [raid.rg.recons.done.notice]: /aggr1/plex0/rg0:
reconstruction completed for v1.16 in 0:13.94

node1> aggr status aggr1
Aggr State      Status          Options
aggr1 online    raid_dp, aggr

Volumes: voll
Plex /aggr1/plex0: online, normal, active
RAID group /aggr1/plex0/rg0: normal

```

Figure 17 Converting existing aggregate to RAID-DP

```

node2> vol status vol2
Volume State      Status           Options
vol2  online       raid4, trad     raidsize=5, create_ucose=on,
                                convert_ucose=on

Containing aggregate: <N/A>
Plex /vol2/plex0: online, normal, active
RAID group /vol2/plex0/rg0: normal

node2> vol options vol2 raidtype raid_dp

WARNING! Continuing with vol options vol2 raidtype raid_dp will result in
having no spare disk available for one or more RAID groups.
Are you sure you want to continue with the operation? y
Mon Mar 20 16:44:52 GMT [raid.config.raidsize.change:notice]: volume vol2:
raidsize is adjusted from 5 to 16 after changing raidtype
Mon Mar 20 16:45:13 GMT [raid.rg.recons.missing:notice]: RAID group
/vol2/plex0/rg0 is missing 1 disk(s).
Mon Mar 20 16:45:13 GMT [raid.rg.recons.info:notice]: Spare disk v0.21 will
be used to reconstruct one missing disk in RAID group /vol2/plex0/rg0.
Mon Mar 20 16:45:13 GMT [raid.rg.recons.start:notice]: /vol2/plex0/rg0:
starting reconstruction, using disk v0.21
Mon Mar 20 16:45:13 GMT [raid.rg.spares.low:warning]: /vol2/plex0/rg0
Mon Mar 20 16:45:25 GMT [raid.rg.recons.done:notice]: /vol2/plex0/rg0:
reconstruction completed for v0.21 in 0:11.94

node2> vol status vol2
Volume State      Status           Options
vol2  online       raid_dp, trad   create_ucose=on,
                                convert_ucose=on

Containing aggregate: <N/A>
Plex /vol2/plex0: online, normal, active
RAID group /vol2/plex0/rg0: normal

```

Figure 18 Converting traditional volume to RAID-DP

Aggregates and traditional volumes may be converted back to RAID4 with the command **[aggr |vol] options name raidtype raid4**. In this case, the conversion is instantaneous, since the old RAID4 row parity construct is still in place as a subsystem in RAID-DP. If a RAID-DP group is converted to RAID4, then each RAID group's second diagonal parity disk is released and put back into the spare disk pool. See Figure 19 and Figure 20.

```

node1> aggr status aggr1
Aggr State      Status          Options
aggr1 online     raid_dp, aggr

Volumes: vol1
Plex /aggr1/plex0: online, normal, active
RAID group /aggr1/plex0/rg0: normal

node1> aggr options aggr1 raidtype raid4
Mon Mar 20 17:06:50 GMT [raid.config.raidsize.change:notice]: Aggregate
aggr1: raidsize is adjusted from 16 to 8 after changing raidtype

node1> aggr status aggr1
Aggr State      Status          Options
aggr1 online     raid4, aggr

Volumes: vol1
Plex /aggr1/plex0: online, normal, active
RAID group /aggr1/plex0/rg0: normal

```

Figure 19 Converting existing RAID-DP aggregate to RAID4

```

node2> vol status vol2
Volume State    Status          Options
vol2  online     raid_dp, trad   create_ucose=on,
                                convert_ucose=on

Containing aggregate: <N/A>
Plex /vol2/plex0: online, normal, active
RAID group /vol2/plex0/rg0: normal

node2> vol options vol2 raidtype raid4
Mon Mar 20 17:11:27 GMT [raid.config.raidsize.change:notice]: Volume vol2:
raidsize is adjusted from 16 to 8 after changing raidtype

node2> vol status vol2
Volume State    Status          Options
vol2  online     raid4, trad     create_ucose=on,
                                convert_ucose=on

Containing aggregate: <N/A>
Plex /vol2/plex0: online, normal, active
RAID group /vol2/plex0/rg0: normal

```

Figure 20 Converting existing RAID-DP traditional volume to RAID4

RAID-DP volume management

From a management and operational point of view, once created or converted to, RAID-DP aggregates and traditional volumes work exactly like their RAID4 counterparts. The same practices and guidelines work for both RAID4- and RAID-DP, so little to no changes are required for standard operational procedures used by IBM System Storage N series administrators. While a storage controller may contain any mix of RAID4 and RAID-DP aggregates or traditional volumes, the commands an administrator would use for management activities on the storage controller are the same.

For instance, if a RAID-DP aggregate or traditional volume requires additional capacity, the command [aggr | vol] add *name* X (where X is the number of disks to add) is run in exactly the same manner as an administrator would use for a RAID4-based storage.

RAID-DP capacity utilization

RAID-DP requires two parity disks per RAID group, and, although this could affect capacity utilization, there are ways to offset the effects of extra disks being used for parity. To reduce or even eliminate the possible impact on capacity utilization from the second parity disk in RAID-DP groups, the first step is to simply use the default RAID-DP group size for the storage platform and disk drive type. Afterwards, create aggregates or traditional volumes in sizes that reflect full increments of the default RAID-DP RAID group size.

Table 2 IBM System Storage N series RAID4 number of disk configuration

Platform	Default	Max
IBM System Storage N3700	7 (6 data + 1 parity)	14 (13 data + 1 parity)
IBM System Storage N5x00	8 (FC) (7 data + 1 parity) 8 (SATA) (7 data + 1 parity)	14 (FC) (13 data + 1 parity) 8 (SATA) (7 data + 1 parity)

Table 3 IBM System Storage N series RAID-DP number of disk configuration

Platform	Default	Max
IBM System Storage N3700	16 (15 data + 1 parity)	28 (27 data + 1 parity)

Platform	Default	Max
IBM System Storage N5x00	16 (FC) (15 data + 1 parity) 14(SATA) (13 data + 1 parity)	28 (FC) (27 data + 1 parity) 16(SATA) (15 data + 1 parity)

Additional information

For information about IBM System Storage N series read information available at:

<http://www-03.ibm.com/servers/storage/nas/index.html>

<http://www.redbooks.ibm.com/redpieces/abstracts/sg247129.html>

The team that wrote this Redpaper

This Redpaper was produced by a team of specialists from around the world working at the International Technical Support Organization, Tucson, Arizona.

Chrisanthy Carlane is an IT Specialist with IBM Information Technology Services in Indonesia. She has five years experience providing enterprise-wide infrastructure implementations, migration and support on IBM Tape and Storage System and IBM @server® xSeries® servers. She has certifications with IBM, Cisco, Microsoft®, Red Hat, and McDATA and holds a bachelor's degree of Economics-Accounting from Tarumanagara University, Jakarta, Indonesia.

Alex Osuna is a Project Leader with the San Jose International Technical Support Organization. He has over 27 years in the IT industry and 22 years of experience in the hardware/software storage area dealing with maintenance, development, early ship programs, education, publishing, performance analysis and technical sales support. He holds 10 certifications from IBM, Microsoft, and Red Hat.

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

Chris Lueth, Network Appliance Technical Sales Manager

Dave Hitz, Network Appliance executive vice president and founder

Notices

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

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 give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 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. 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 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 IBM product and use of those Web sites is at your own risk.

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

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.

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 illustrates 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. 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 IBM's application programming interfaces.

© Copyright International Business Machines Corporation 2006. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Send us your comments in one of the following ways:


- ▶ Use the online **Contact us** review redbook form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbook@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099, 2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.



Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server®

Redbooks (logo) ™

xSeries®

IBM®

Redbooks™

System Storage™

The following terms are trademarks of other companies:

Data ONTAP, SyncMirror, FilerView, Network Appliance, The Network Appliance logo, the bolt design, Camera-to-Viewer, Center-to-Edge, ContentDirector, ContentFabric, NetApp Availability Assurance, NetApp ProTech Expert, NOW, NOW NetApp on the Web, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, Smart SAN, The evolution of storage, Virtual File Manager, and Web Filer are trademarks of Network Appliance, Inc. in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

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