



Common RAID Disk Data Format Specification

Version 2.0 Revision 19

Publication of this *Working Draft* for review and comment has been approved by the DDF TWG. This draft represents a “best effort” attempt by the DDF TWG to reach preliminary consensus, and it may be updated, replaced, or made obsolete at any time. This document should not be used as reference material or cited as other than a “work in progress.” Suggestion for revision should be directed to <http://www.snia.org/feedback/>.

Working Draft

November 21, 2008

Revision History

Revision	Date	Sections	Originator:	Comments
1.2	7/28/2006		Arnold Jones	Officially published as SNIA Technical Position.
1.20.01	8/16/2006		Bill Dawkins	Development draft of the next version of the DDF Specification
1.20.02	2/22/2007		Bill Dawkins	Added content to support large block size drives
1.20.03	03/02/2007		Ramamurthy Krithivas	Added RAID 5 Rotate after N Stripes equations, and RAID 6 PQ ordered equation.
1.20.04	3/6/2007		Bill Dawkins	Modified content to support large block size drives.
1.20.05	5/3/2007	Various	Bill Dawkins	Updated draft to include all 1.2 errata. Added variable entries to BBM_Log section. Added enhanced definition of VD states to account for MDF RAID. Added path information for SATA drives in PDE structures.
1.20.06	06/21/2007	Various	Ramamurthy Krithivas	First Draft of the MDF RAID Algorithm.
1.20.07	07/30/2007	4.2.25	Ramamurthy Krithivas	Updated 4.2.25 based on TWG review of 06/26/2007.
1.20.11	11/21/2007		Bill Dawkins	Draft for IP review
Version 2.0 rev 17	10/8/2008	Various	Bill Dawkins	Draft candidate for Technical Council submission.
Version 2.0 rev 18	11/18/2008	Various	Bill Dawkins	Updated draft candidate for Technical Council submission.
Version 2.0 rev 19	11/21/2008	Various	Bill Dawkins	Approved draft submitted to Technical Council

Suggestion for changes or modifications to this document should be submitted at <http://www.snia.org/feedback/>.

Typographical Conventions

The key words “**MUST**”, “**MUST NOT**”, “**REQUIRED**”, “**SHALL**”, “**SHALL NOT**”, “**SHOULD**”, “**SHOULD NOT**”, “**RECOMMENDED**”, “**MAY**”, and “**OPTIONAL**” in this document are to be interpreted as described in RFC2119 [<http://www.ietf.org/rfc/rfc2119.txt>].

Usage

The SNIA hereby grants permission for individuals to use this document for personal use only, and for corporations and other business entities to use this document for internal use only (including internal copying, distribution, and display) provided that:

- 1) Any text, diagram, chart, table or definition reproduced must be reproduced in its entirety with no alteration; and
- 2) Any document, printed or electronic, in which material from this document (or any portion hereof) is reproduced must acknowledge the SNIA copyright on that material, and must credit the SNIA for granting permission for its reuse.

Other than as explicitly provided above, you may not make any commercial use of this document, sell any or this entire document, or distribute this document to third parties. All rights not explicitly granted are expressly reserved to SNIA.

Permission to use this document for purposes other than those enumerated above may be requested by e-mailing td@snia.org. Please include the identity of the requesting individual and/or company and a brief description of the purpose, nature, and scope of the requested use.

Copyright © 2005 – 2008, Storage Networking Industry Association.

TABLE OF CONTENTS

1	INTRODUCTION.....	8
2	OVERVIEW.....	9
2.1	Purpose.....	9
2.2	Design Considerations	10
2.2.1	Location.....	10
2.2.2	Locality.....	10
2.2.3	DDF Structure Size	10
2.2.4	DDF Structure Contents	10
2.2.5	DDF Structure Redundancy	11
3	DEFINITIONS.....	12
3.1	RAID Terms.....	12
3.1.1	Virtual Disk (VD).....	12
3.1.2	Basic Virtual Disk (BVD).....	12
3.1.3	Secondary Virtual Disk (SVD)	12
3.1.4	Disk Grouping.....	12
3.1.5	Foreign configuration.....	12
3.1.6	Legacy or Pass-through Disk	12
4	RAID LEVELS AND RAID LEVEL QUALIFIERS	13
4.1	Primary RAID Level	13
4.2	RAID Level Qualifier	13
4.2.1	RAID-0 Simple Striping (PRL=00, RLQ=00).....	16
4.2.2	RAID-1 Simple Mirroring (PRL=01, RLQ=00).....	18
4.2.3	RAID-1 Multi Mirroring (PRL=01, RLQ=01)	19
4.2.4	RAID-3 Non-Rotating Parity 0 (PRL=03, RLQ=00).....	20
4.2.5	RAID-3 Non-Rotating Parity N (PRL=03, RLQ=01)	22
4.2.6	RAID-4 Non-Rotating Parity 0 (PRL=04, RLQ=00).....	24
4.2.7	RAID-4 Non-Rotating Parity N (PRL=04, RLQ=01).....	26
4.2.8	RAID-5 Rotating Parity 0 with Data Restart (PRL=05, RLQ=00)	28
4.2.9	RAID-5 Rotating Parity N with Data Restart (PRL=05, RLQ=02)	31
4.2.10	RAID-5 Rotating Parity N with Data Continuation (PRL=05, RLQ=03)	33
4.2.11	RAID-5E Rotating Parity 0 with Data Restart (PRL=15, RLQ=00).....	35
4.2.12	RAID-5E Rotating Parity N with Data Restart (PRL=15, RLQ=02)	38
4.2.13	RAID-5E Rotating Parity N with Data Continuation (PRL=15, RLQ=03) ..	40
4.2.14	RAID-5EE Rotating Parity 0 with Data Restart (PRL=25, RLQ=00)	42
4.2.15	RAID-5EE Rotating Parity N with Data Restart (PRL=25, RLQ=02).....	45
4.2.16	RAID-5EE Rotating Parity N with Data Continuation (PRL=25, RLQ=03) ..	48
4.2.17	RAID-5R Rotating Parity 0 after R Stripes with Data Restart (PRL=35, RLQ=00)50	
4.2.18	RAID-5R Rotating Parity N after R Stripes with Data Restart (PRL=35, RLQ=02)53	

4.2.19	<i>RAID-5R Rotating Parity N after R Stripes with Data Continuation (PRL=35, RLQ=03)</i>	55
4.2.20	<i>RAID-1E Integrated Adjacent Stripe Mirroring (PRL= 11, RLQ=00)</i>	57
4.2.21	<i>RAID-1E Integrated Offset Stripe Mirroring (PRL=11, RLQ=01)</i>	59
4.2.22	<i>RAID-6 Rotating Parity 0 with Data Restart (PRL=06, RLQ=01)</i>	61
4.2.22.1	Parity Re-computation on Block Update	64
4.2.22.2	Galois Field Operations	64
4.2.23	<i>RAID-6 Rotating Parity N with Data Restart (PRL=06, RLQ=02)</i>	68
4.2.23.1	Parity Re-computation on Block Update	70
4.2.23.2	Galois Field Operations	70
4.2.24	<i>RAID 6 Rotating Parity N with Data Continuation (PRL=06, RLQ=03)</i>	71
4.2.24.1	Parity Re-computation on Block Update	73
4.2.24.2	Galois Field Operations	73
4.2.25	<i>Multi Disk Failure RAID Rotating Parity 0 with Data Restart (PRL=07, RLQ=00)</i>	74
4.2.25.1	Galois Field Generation	76
4.2.25.2	Constant Matrix Generation	77
4.2.25.3	Galois Parity Computation	79
4.2.26	<i>Multi Disk Failure RAID Rotating Parity N with Data Restart (PRL=07, RLQ=02)</i>	80
4.2.27	<i>Multi Disk Failure RAID Rotating Parity N with Data Continuation (PRL=07, RLQ=03)</i>	82
4.3	Secondary RAID Level	84
4.3.1	<i>Striped Secondary RAID Level (SRL=00)</i>	85
4.3.2	<i>Mirrored Secondary RAID Level (SRL=01)</i>	87
4.3.3	<i>Concatenated Secondary RAID Level (SLR=02)</i>	88
4.3.4	<i>Spanned Secondary RAID Level (SRL=03)</i>	90
5	DDF STRUCTURE	93
5.1	DDF Structure Overview	93
5.2	Byte Ordering	94
5.3	Signatures, Timestamps and CRCs	97
5.4	GUIDs	97
5.4.1	<i>Controller GUID</i>	97
5.4.2	<i>Physical Disk GUID</i>	98
5.4.3	<i>Virtual Disk GUID</i>	99
5.4.4	<i>DDF Header GUID</i>	99
5.5	DDF Header	99
5.6	Controller Data	105
5.7	Physical Disk Records	106
5.7.1	<i>Physical Disk Entries</i>	107
5.8	Virtual Disk Records	110
5.8.1	<i>Virtual Disk Entries</i>	111
5.9	Configuration Records	114
5.9.1	<i>Virtual Disk Configuration Record</i>	114
5.9.2	<i>Vendor Unique Configuration Record</i>	119
5.9.3	<i>Spare Assignment Record</i>	120
5.9.3.1	Spare Assignment Entry	121
5.9.3.2	Spare Types and Spare Assignment Entry Details	122

5.10 For committable dedicated spares (Spare Type: Bit1 = 0), once the disk joins a VD, it becomes a permanent member of the VD and the Spare Assignment Record on the spare associated with the VD MUST be deleted.Physical Disk Data
123

5.11 Bad Block Management Log	124
5.11.1 <i>Mapped/Marked Block Entry</i>	<i>124</i>
5.12 Diagnostic Space	125
5.13 Vendor Specific Logs.....	125

1 Introduction

In today's IT environments, there are several reasons why system administrators would wish to change the internal RAID solutions they are using. For example, many servers are shipped with a RAID solution implemented on the motherboard (ROMB). ROMB solutions allow RAID formats to be applied to the disks internal to the server. As the server's data set grows, the administrator often finds s/he needs to move to a larger direct attached storage (DAS) solution with external JBODs. The system administrator would like to move the internal disks and their data to the DAS system's external JBODs. One method of migration is to backup a RAID group, transfer the disks to the new storage system, reconfigure the disks as a new RAID group behind the new RAID controller, and restore the data from the backup device. This time consuming procedure also carries some risk of data loss. A better method would be to move the disks with data-in-place from one RAID implementation to another. Unfortunately, the different methods for storing configuration information prohibit data-in-place migration between systems from different storage vendors.

The SNIA Common RAID Disk Data Format Technical Working Group was chartered define a standard data structure describing how data is formatted across the disks in a RAID group. This specification defines the Common RAID Disk Data Format (DDF) structure. The DDF structure allows a basic level of interoperability between different suppliers of RAID technology. The Common RAID DDF structure benefits storage users by enabling data-in-place migration among systems from different vendors.

Part of the specification defines how data is distributed for many basic RAID levels. This is necessary to precisely document how data is formatted for RAID levels indicated by the DDF structure. The DDF TWG recognizes that the formats described do not represent all methods for implementing the defined RAID levels. The SNIA does not imply that specification formats represent a preferred RAID implementation. Reviewers of this specification are encouraged to suggest alternate RAID level formats for inclusion into future revisions of the specification.

The DDF data structure also allows RAID groups with vendor unique RAID formats. While vendor unique RAID formats prohibit data-in-place migration between vendors, the Common RAID DDF will be a benefit in these situations. At a minimum, when a RAID group containing a unique format is moved to a different RAID solution that does not support the format, the new system will still be able to read the DDF structure. It can identify the RAID groups containing the unique format and notify the system administrator that these RAID groups contain valid data that is not accessible by the current storage system. Potential data loss is prevented because the new RAID system will not overwrite the data without administrator confirmation.

The document is divided into the following sections:

- Section 2 is the overview;
- Section 3 describes the definitions used in this specification;
- Section 4 describes the RAID levels defined in this specification;
- Section 5 details the DDF structure.

2 Overview

2.1 Purpose

This document provides requirements for the RAID configuration Disk Data Format (DDF) structure stored on physical disks attached to RAID controllers. Configuration on Disk (COD) and Metadata are also commonly used terms for the same type of data structure. This DDF structure allows storing RAID configuration information on physical disks by different vendor implementations in a common format so the user data on physical disks is accessible independent of the RAID controller being used. Controllers are not required to store this information in the same format in their internal memory.

In the terminology of the SNIA Shared Storage Model (http://www.snia.org/tech_activities/shared_storage_model), the technical scope of the DDF is limited to the interface between a *block aggregation* implementation and *storage devices*. The DDF is stored as data structures on the *storage devices* (see Figure 1).

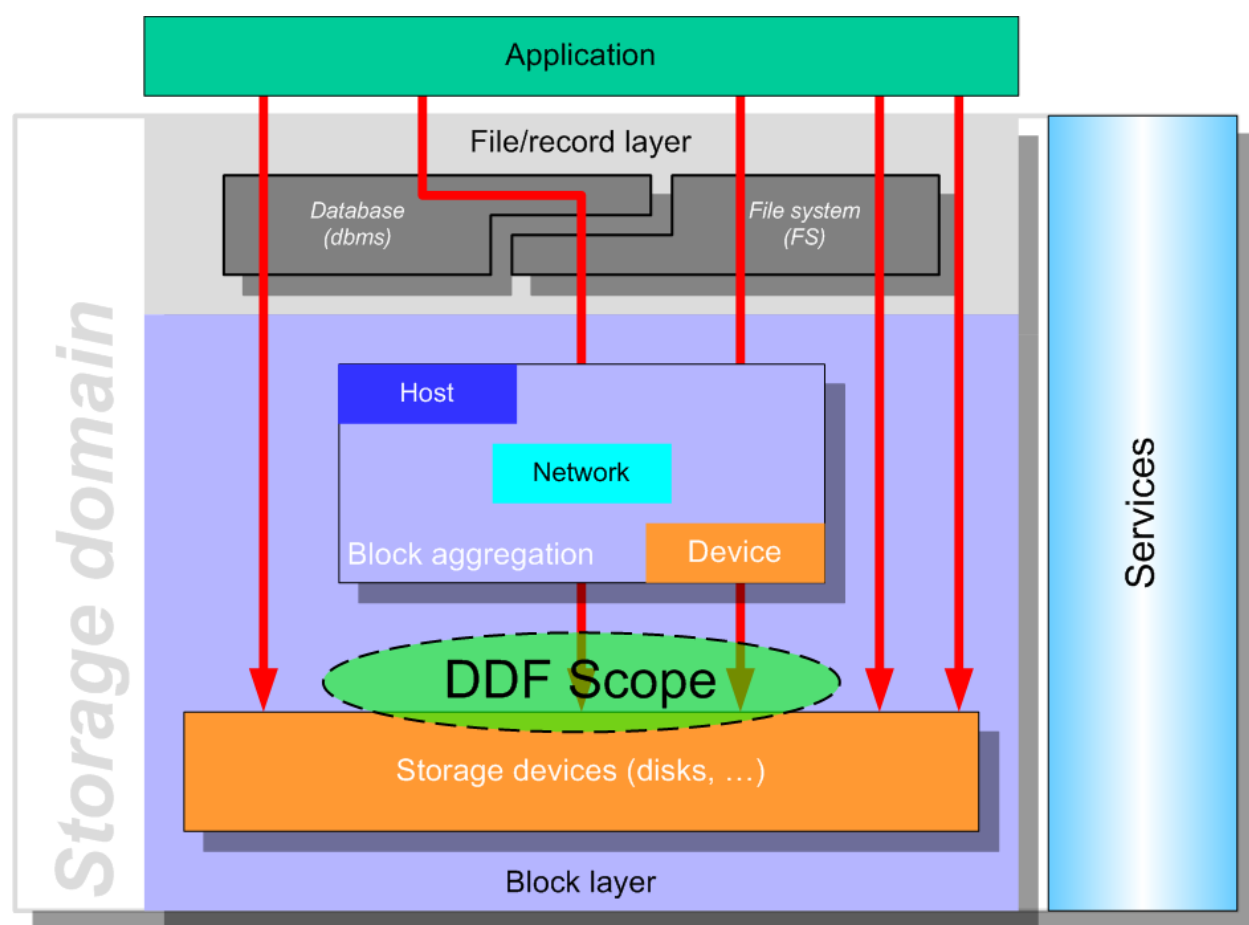


Figure 1: DDF Technical Scope

2.2 Design Considerations

Location, locality, size and contents are major considerations for the DDF structure. Details on the contents and format of the DDF structure can be found in Section 5.

2.2.1 Location

The Anchor Header (see Section 5.5) for the DDF structure MUST be stored at the last logical block returned by either the ATA Identify Device or SCSI Read Capacity commands depending on the type of physical disk.

The DDF structure SHOULD be stored at the end of the physical disk next to the anchor header. Storing the DDF structure at the end of the physical disk allows the possibility of converting a single non-RAID physical disk to a RAID 1 configuration without shifting user data. Similarly, data on a member of a RAID 1 configuration with the DDF structure at the end can also be accessed using a non-RAID controller.

2.2.2 Locality

Locality is defined as the scope of the DDF structure. One approach is to store configuration information about all RAID groups attached to a controller on every physical disk connected to the controller. The second approach is to store configuration information of a RAID group (or virtual disk) only on the physical disks belonging to the participating in a RAID group. In other words, does the DDF structure on one RAID group have information about other RAID groups attached to the same controller? This plays a role when a user wants to move a RAID group from one controller to another while keeping the RAID group intact and without causing ghost images on either controller. If the DDF structure on one RAID group contains no information about the other RAID groups and if an entire RAID group disappears due to power or cabling problems, the user should be notified about the missing RAID group. Configuration information about all RAID groups may be stored on NVRAM on the controller and provide a notification for the user. However, in case of a failed or replacement controller, the information on the RAID groups is not available.

The middle ground, chosen for the DDF structure, is to store the complete configuration information of a RAID group on each physical disk participating in the RAID group and to store a minimal amount of information about other the RAID groups and physical disks attached to the controller. This allows a controller to provide notification to the user about missing RAID groups or physical disks when the controller does not have complete information about the configuration.

2.2.3 DDF Structure Size

A large DDF structure size provides room for expansion in the future and still uses a negligible amount of storage. It is tempting to use a fixed large DDF structure size; however, low end solutions may not have the memory space to process large tables.

The DDF structure size is not fixed and depends on solution needs. This is done by using flexible structures where size is a function of the number of physical and virtual disks. A fixed space SHOULD be reserved on physical disks for the DDF structure to accommodate the largest DDF structure size for migration of configurations between different types and classes of solutions. Details on DDF structure size can be found in Section 5.

2.2.4 DDF Structure Contents

The DDF structure contains information about partitioning, RAID level, and cache parameters for each virtual disk defined. RAID group state, physical disk location information, and controller settings are among other information included in the DDF structure. The DDF structure contents are defined in detail in Section 5.

2.2.5 DDF Structure Redundancy

DDF structure redundancy allows recovery after configuration structure corruption or loss. Support for redundancy increases solution complexity and is considered OPTIONAL.

3 Definitions

Whenever possible, this specification uses the definitions for storage terminology provided by SNIA's "A Dictionary of Storage Networking Terminology." The dictionary can be found at <http://www.snia.org/education/dictionary>. This section defines terms that do not have entries in the SNIA dictionary. It also defines terms that use definitions that differ from the definition listed in the SNIA dictionary.

3.1 RAID Terms

3.1.1 Virtual Disk (VD)

A virtual disk is the object presented to the host level for user data storage. At least one physical disk is associated with a VD.

3.1.2 Basic Virtual Disk (BVD)

A basic virtual disk is a VD configured using only non-hybrid RAID levels like RAID-0, RAID5 or RAID5E. Its elements are physical disks.

3.1.3 Secondary Virtual Disk (SVD)

A secondary virtual disk is a VD configured using hybrid RAID levels like RAID10 or RAID50. Its elements are BVDs.

3.1.4 Disk Grouping

A number of physical disks can be combined into a disk group. The primary characteristic of a disk group is that all VDs created on the physical disks cannot extend to other physical disks that are not part of the group. Disk Grouping, when enforced, ensures that the contiguous address space of a VD does not extend beyond a disk group.

3.1.5 Foreign configuration

A configuration moved from one controller to another controller is considered a foreign configuration on the new controller unless new controller imports the configuration. Whenever a foreign configuration is detected by a controller, the Foreign_Flag MAY be set in the DDF header on the physical disks in the foreign configuration. Details of the Foreign_Flag are found in Section 5.5.

3.1.6 Legacy or Pass-through Disk

Legacy (pass-through) physical disks are attached to a RAID controller and operate as if they were attached to a non-RAID controller. No DDF structure is stored on these physical disks. This feature is primarily targeted for users moving physical disks containing data from non-RAID controllers to RAID controllers.

4 RAID Levels and RAID Level Qualifiers

This section lists the RAID types and qualifiers for use in following fields used in the Configuration Record (Section 5.9):

- Primary RAID Level
- RAID Level Qualifier
- Secondary RAID Level

4.1 Primary RAID Level

Table 1 lists values used in the Primary_RAID_Level field of the Virtual Disk Configuration Record (Section 5.9.1) and the definitions of these values. The Primary_RAID_Level field MUST use the values listed in Table 1. The table defines the standard RAID levels, such as RAID 0, 1, 3, 5, etc. and some proprietary RAID types. Non-RAID types such as JBOD and concatenation are also included for completeness.

Table 1: Primary RAID Levels

<u>Name</u>	<u>PRL Byte</u>	<u>Description</u>
RAID-0	0x00	Striped array with no parity
RAID-1	0x01	Mirrored array
RAID-3	0x03	Striped array with non-rotating parity, optimized for long, single-threaded transfers
RAID-4	0x04	Striped array with non-rotating parity, optimized for short, multi-threaded transfers
RAID-5	0x05	Striped array with rotating parity, optimized for short, multi-threaded transfers
RAID-6	0x06	Similar to RAID-5, but with dual rotating parity physical disks, tolerating two physical disk failures
MDF RAID	0x07	Multi disk Failure RAID. Similar to RAID-6, but supporting more than two physical disk failures
RAID-1E	0x11	>2 disk RAID-1, similar to RAID-10 but with striping integrated into array
Single Disk	0x0F	Single, non-arrayed disk
Concatenation	0x1F	Physical disks combined head to tail
RAID-5E	0x15	RAID-5 with hot space at end each extent
RAID-5EE	0x25	RAID-5 with hot space integrated into each extent
RAID-5R	0x35	RAID-5 with parity rotated after a configured number of stripes

4.2 RAID Level Qualifier

This section defines RAID Level Qualifiers (RLQ) for each Primary RAID Level as defined earlier. The RLQ field MUST be ignored for JBOD and Concatenations (PRL=0F and 1F). Table 2 gives brief descriptions of the RLQs defined for each Primary RAID Level described in this specification. The following subsections describe the data formats defined by the PRLs and the associated RLQs in detail.

Table 2: RAID Level Qualifiers

<u>Name</u>	<u>PRL Byte</u>	<u>RLQ Byte</u>	<u>Description</u>
RAID-0	0x00	0x00	No RLQs for this PRL.
RAID-1 Simple Mirroring	0x01	0x00	Mirroring across two extents.
RAID-1 Multi Mirroring	0x01	0x01	Triple mirroring across three extents.
RAID-3 Non-Rotating Parity 0	0x03	0x00	Striped array with parity stored on the first extent.
RAID-3 Non-Rotating Parity N	0x03	0x01	Striped array with parity stored on the last extent.
RAID-4 Non-Rotating Parity 0	0x04	0x00	Striped array with parity stored on the first extent.
RAID-4 Non-Rotating Parity N	0x04	0x01	Striped array with parity stored on the last extent.
RAID-5 Rotating Parity 0 with Data Restart	0x05	0x00	Striped array with rotating parity striped diagonally from left to right. The first data strip of each stripe starts on the first extent not occupied by parity.
RAID-5 Rotating Parity N with Data Restart	0x05	0x02	Striped array with rotating parity striped diagonally from right to left. The first data strip of each stripe starts on the first extent not occupied by parity.
RAID-5 Rotating Parity N with Data Continuation	0x05	0x03	Striped array with rotating parity striped diagonally from right to left. The first data strip of each stripe starts directly below the parity strip of the previous stripe.
RAID-6 Rotating Parity 0 with Data Restart	0x06	0x00	Dual rotating parity striped diagonally from left to right. The first data strip of each stripe starts on the first extent not occupied by parity.
RAID-6 Rotating Parity N with Data Restart	0x06	0x02	Dual rotating parity striped diagonally from right to left. The first data strip of each stripe starts on the first extent not occupied by parity.
RAID-6 Rotating Parity N with Data Continuation	0x06	0x03	Dual rotating parity striped diagonally from right to left. The first data strip of each stripe starts directly below the second parity strip of the previous stripe.
MDF RAID Rotating Parity 0 with Data Restart	0x07	0x00	Multi Disk Failure RAID. Multiple rotating parity striped diagonally from left to right. The first data strip of each stripe starts on the first extent not occupied by parity.
MDF RAID Rotating Parity N with Data Restart	0x07	0x02	Multi Disk Failure RAID. Multiple rotating parity striped diagonally from right to left. The first data strip of each stripe starts on the first extent not occupied by parity.
MDF RAID Rotating Parity N with Data Continuation	0x07	0x03	Multi Disk Failure RAID. Multiple rotating parity striped diagonally from right to left. The first data strip of each stripe starts on directly below the last parity of the previous strip.
RAID-1E Integrated Adjacent Stripe Mirroring	0x11	0x00	Each data strip is mirrored on the next extent. For arrays with an odd number of extents, one strip of each stripe is mirrored on the next or previous stripe.
RAID-1E Integrated Offset Stripe Mirroring	0x11	0x01	Each data strip is mirrored on the next stripe of the next extent.

Single Disk	0x0F	0x00	No RLQs for this PRL.
Concatenation	0x1F	0x00	No RLQs for this PRL.
RAID-5E Rotating Parity 0 with Data Restart	0x15	0x00	Same as PRL=05, RLQ=00 with hot space at the end of each extent.
RAID-5E Rotating Parity N with Data Restart	0x15	0x02	Same as PRL=05, RLQ=02 with hot space at the end of each extent.
RAID-5E Rotating Parity N with Data Continuation	0x15	0x03	Same as PRL=05, RLQ=03 with hot space at the end of each extent.
RAID-5EE Rotating Parity 0 with Data Restart	0x25	0x00	Same as PRL=05, RLQ=00 with hot space integrated in to each extent. The hot space of each stripe directly follows the parity block of that stripe.
RAID-5EE Rotating Parity N with Data Restart	0x25	0x02	Same as PRL=05, RLQ=02 with hot space integrated in to each extent. The hot space of each stripe directly precedes the parity block of that stripe.
RAID-5EE Rotating Parity N with Data Continuation	0x25	0x03	Same as PRL=05, RLQ=03 with hot space integrated in to each extent. The hot space of each stripe directly precedes the parity block of that stripe.
RAID-5 Rotating Parity 0 after R Stripes with Data Restart	0x35	0x00	Same as PRL=05, RLQ=00 with parity rotated after R stripes instead of every stripe.
RAID-5 Rotating Parity N after R Stripes with Data Restart	0x35	0x02	Same as PRL=05, RLQ=02 with parity rotated after R stripes instead of every stripe.
RAID-5 Rotating Parity N after R Stripes with Data Continuation	0x35	0x03	Same as PRL=05, RLQ=03 with parity rotated after R stripes instead of every stripe.

4.2.1 RAID-0 Simple Striping (PRL=00, RLQ=00)

Figure 2 shows an example of simple striping (RAID-0).

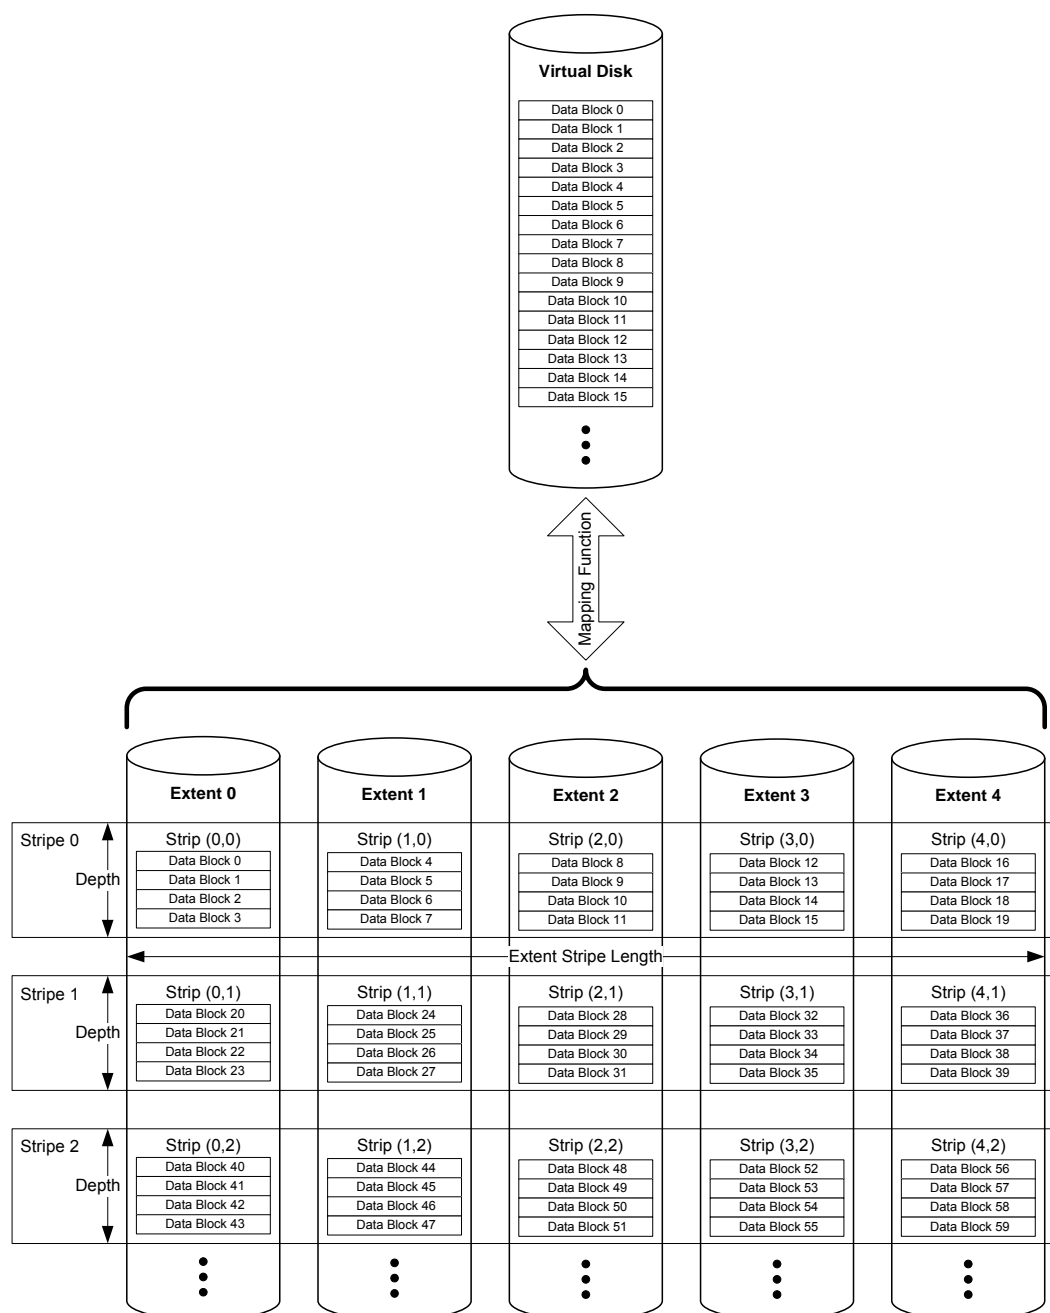


Figure 2: Simple Striping (PRL=00, RLQ=00) Example

The standard SNIA dictionary definitions for stripe, strip, stripe depth, and extent are used by this example and following examples. For a Basic Virtual Disk (BVD), as defined by this specification, an extent MUST be a contiguous area of a physical disk. A BVD's extents MUST be of equal size but are not required to reside in the same location of each physical disk.

The example introduces the concept of *extent* and *stripe* indices for a strip. *strip* (i, j) represents the strip located on extent i in stripe j . The *data block* index k represents the offset of a given data block from the beginning of a strip. To represent a specific block in a specific extent the following notation is used:

extent_block (k, i, j).

To refer to a specific block in a virtual disk, the following notation is used:

virtual_block (x),

where x is the offset from the beginning of the VD.

Table 3 summarizes the indices and constants used in the notation along with the any restrictions on legal values.

Table 3: Indices and Constants for Simple Striping

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD	0	$N-1$
j	Index of a stripe in a VD	0	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/N)$
k	Offset of a data block from the beginning of a strip	0	$L-1$
L	Size of a strip in blocks	N/A (a fixed value)	N/A (a fixed value)
M	Number of blocks in a VD. M MUST be evenly divisible by N .	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD	N/A (a fixed value)	N/A (a fixed value)
x	Offset of a data block from the beginning of a VD	0	$M-1$

For RAID-0 (PRL=00, RLQ=00), the first strip of the virtual disk MUST be strip (0,0). The allocation of data MUST adhere to the following formula:

Eq. 1

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), \text{MOD}(\text{FLOOR}(x/L), N), \text{FLOOR}(\text{FLOOR}(x/L)/N)).$$

4.2.2 RAID-1 Simple Mirroring (PRL=01, RLQ=00)

A VD with PRL=01 and RLQ=00 MUST have two and only two extents. Each extent MUST be equal to the size of the VD. Each block of the VD, virtual_block(x), MUST be duplicated on both extents at the same offset

Figure 3 gives an example of simple mirroring.




Figure 3: Simple Mirroring (PRL=01, RLQ=00) Example

4.2.3 RAID-1 Multi Mirroring (PRL=01, RLQ=01)

Multi Mirroring (PRL=01, RLQ=01) is a triple mirror. Data **MUST** be triple copied on three extents. Each virtual_block(x) **MUST** be duplicated on each extent in the VD. Figure 4 gives an example of multi mirroring.

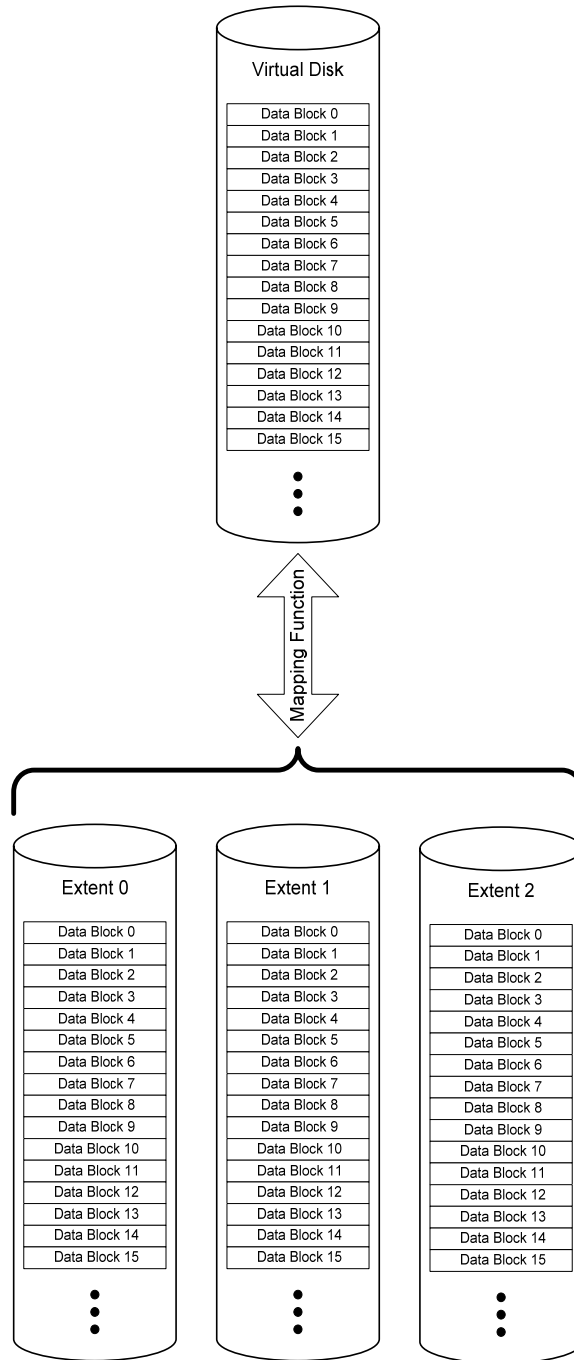


Figure 4: Multi Mirroring (PRL=01, RLQ=01) Example

4.2.4 RAID-3 Non-Rotating Parity 0 (PRL=03, RLQ=00)

Figure 5 gives an example of RAID-3 with parity contained on the first extent or Non-Rotating Parity 0. Table 4 defines the indices and constants used in the description of RAID-3.

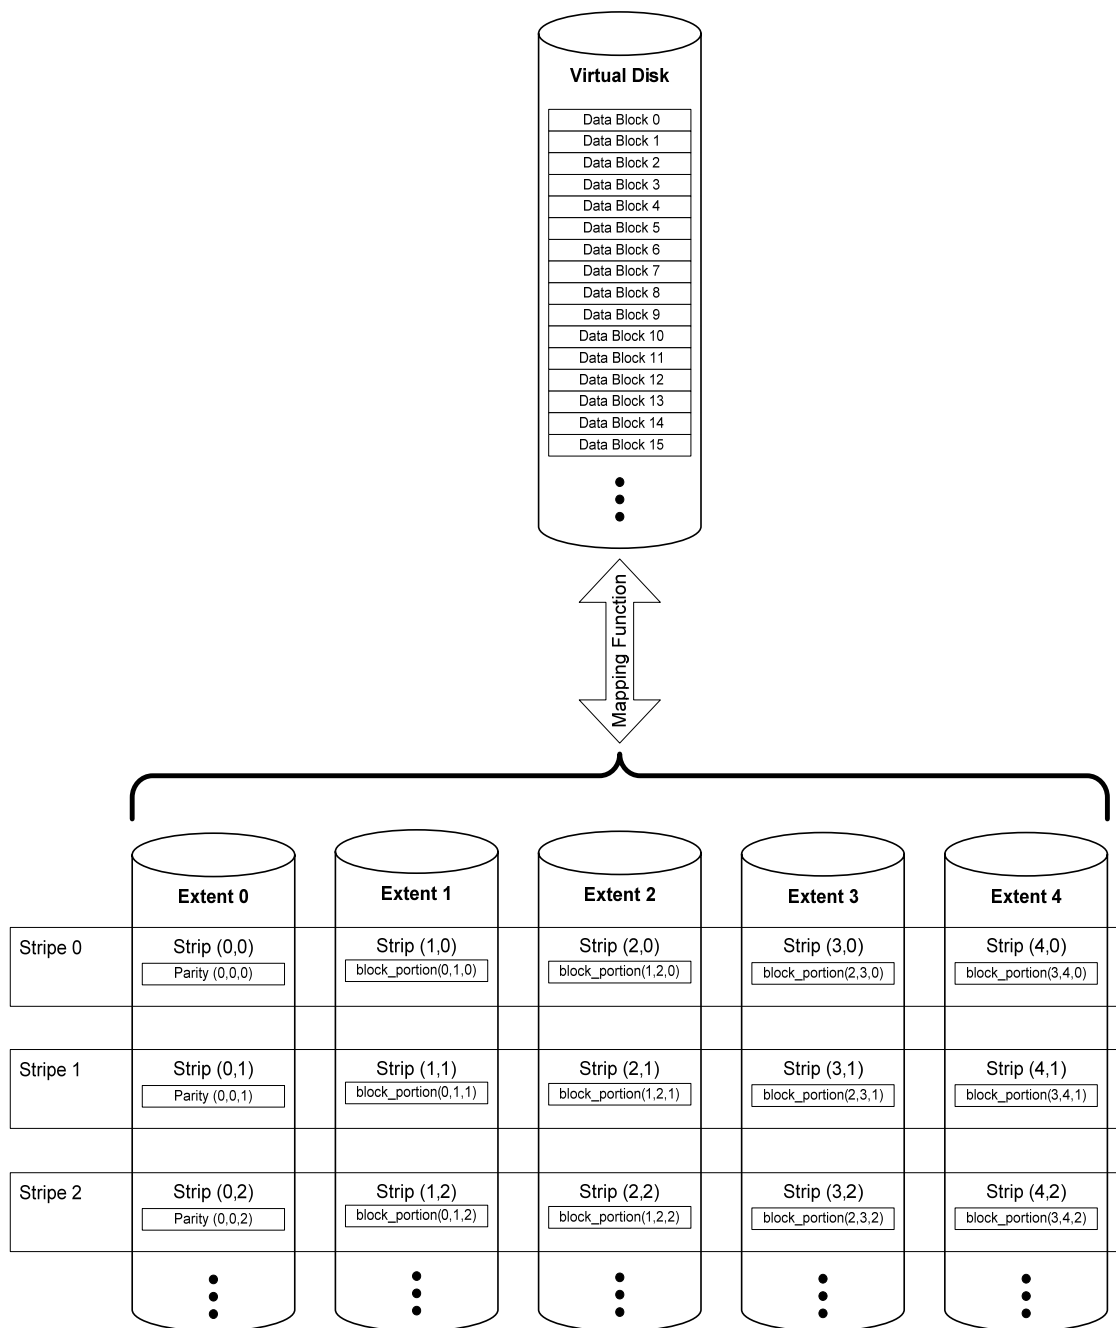


Figure 5: RAID-3 Non-Rotating Parity 0 (PRL=03, RLQ=00) Example

Table 4: Indices and Constants for RAID-3 Non-Rotating Parity

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD	0	$N-1$
j	Index of a stripe in a VD	0	$M-1$
p	Index of a data block portion	0	$N-1$
M	Number of data blocks in a VD	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD	N/A (a fixed value)	N/A (a fixed value)
x	Offset of a data block from the beginning of a VD	0	$M-1$

In a RAID-3 VD with N extents, a virtual data block MUST be segmented into $N-1$ block portions. The notation for data block portions is:

block_portion (p, i, j)

The allocation of data blocks in a RAID-3 Non-Rotating Parity 0 VD MUST adhere to the following formula:

Eq. 2

$$\text{virtual_block}(x) = \parallel_{p=0}^{N-2} \text{block_portion}(p, p+1, x),$$

where \parallel represents the concatenation operator.

Parity blocks MUST reside on extent 0. The values of the parity blocks must adhere to the following formula:

Eq. 3

$$\text{parity_block}(0, 0, x) = \bigoplus_{p=0}^{N-2} \text{block_portion}(p, p+1, x).$$

The operator \oplus refers to bit-wise XOR of the operands.

4.2.5 RAID-3 Non-Rotating Parity N (PRL=03, RLQ=01)

Figure 6 gives an example of a RAID-3 Non-Rotating Parity N VD. The indices and constants defined in Table 4 are valid for this type of VD.

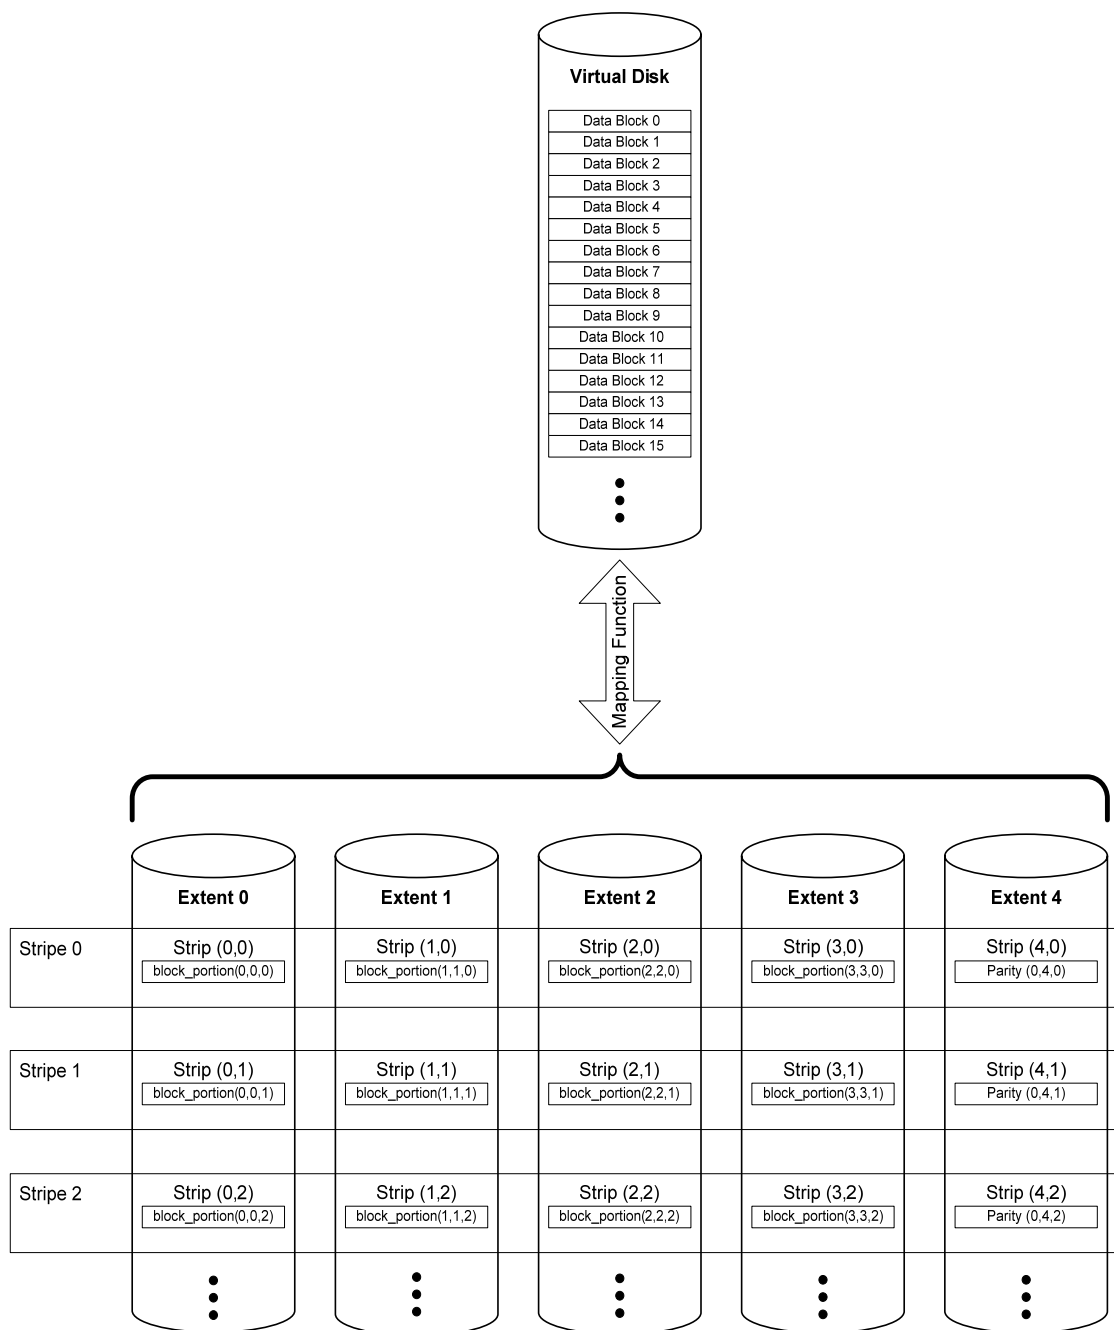


Figure 6: RAID-3 Non-Rotating Parity N (PRL=03, RLQ=01) Example

The allocation of data blocks in a RAID-3 Non-Rotating Parity N VD MUST adhere to the following formula:

Eq. 4

$$\text{virtual_block}(x) = \parallel_{p=0}^{N-2} \text{block_portion}(p, p, x).$$

Parity Blocks MUST reside on extent N . The allocation of parity blocks MUST adhere to the following formula:

Eq. 5

$$\text{parity_block}(0, N, x) = \bigoplus_{p=0}^{N-2} \text{block_portion}(p, p, x).$$

4.2.6 RAID-4 Non-Rotating Parity 0 (PRL=04, RLQ=00)

Figure 7 gives an example of RAID-4 Non-Rotating Parity 0. Table 5 defines the indices and constants used in the description of RAID-4.

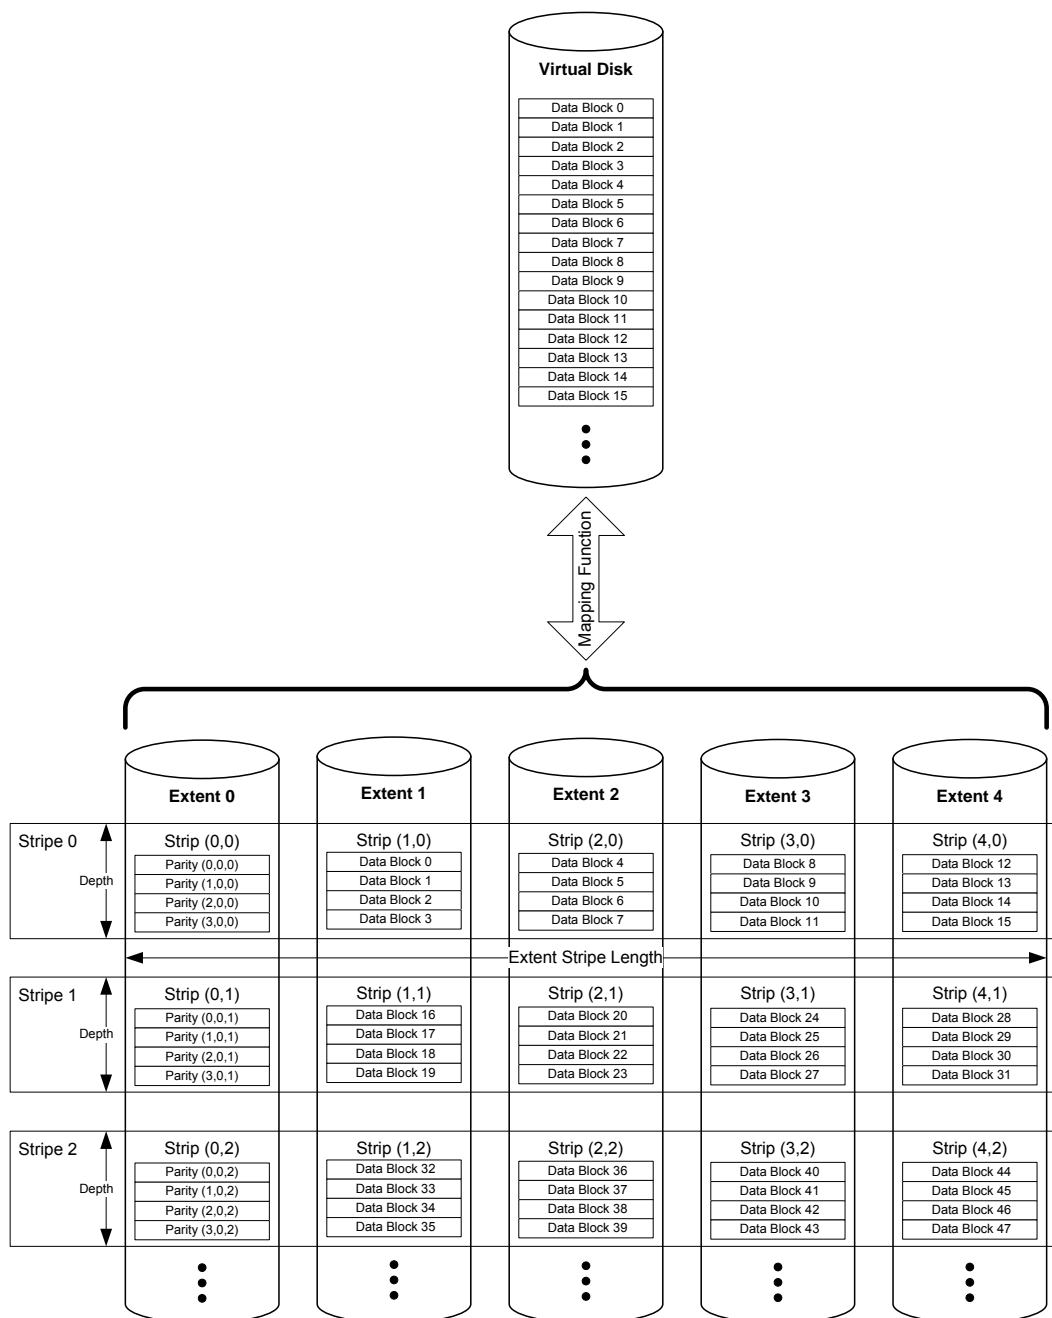


Figure 7: Non-Rotating Parity 0 (PRL=04, PRL=00) Example

Table 5: Indices and Constants for RAID 4 Non-Rotating Parity

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD	0	$N-1$
j	Index of a stripe in a VD	0	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/N-1)$
k	Offset of a block from the beginning of a strip	0	$L-1$
L	Size of a strip in blocks	N/A (a fixed value)	N/A (a fixed value)
M	Number of data blocks in a VD. M MUST be evenly divisible by $N-1$.	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD	N/A (a fixed value)	N/A (a fixed value)
x	Offset of a data block from the beginning of a VD	0	$M-1$

The allocation of data blocks in a Non-Rotating Parity 0 VD MUST adhere to the following equation:

Eq. 6

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), \text{MOD}(\text{FLOOR}(x/L), N-1) + 1, \text{FLOOR}(\text{FLOOR}(x/L)/N-1)).$$

A parity block contains the parity calculated for $N-1$ data blocks. The following notation is used to represent a parity block:

$$\text{parity_block}(k, i, j).$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 7

$$\text{parity_block}(k, 0, j) = \bigoplus_{i=1}^{N-1} \text{extent_block}(k, i, j).$$

For Non-Rotating Parity 0, all parity blocks MUST reside on extent 0. Thus, i MUST equal zero for all parity blocks.

4.2.7 RAID-4 Non-Rotating Parity N (PRL=04, RLQ=01)

Figure 8 gives an example on Non-Rotating Parity N (a.k.a. RAID-4). Non-Rotating Parity N differs from Non-Rotating Parity 0 in that the parity is stored in the last extent of a VD. The indices and constants of Table 5 are valid for Non-Rotating Parity N.

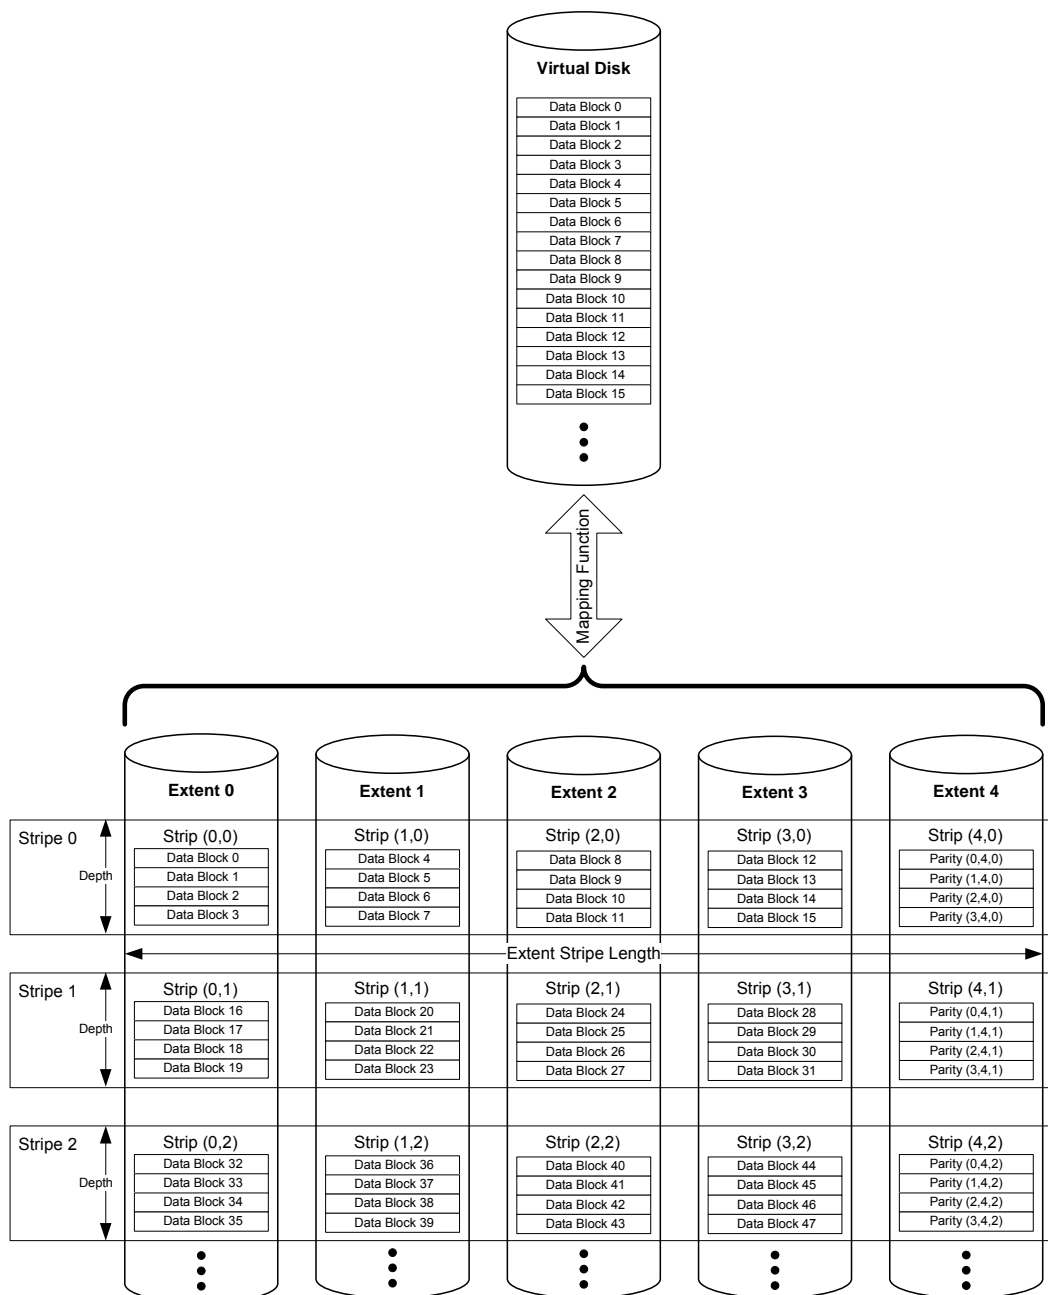


Figure 8: Non-Rotating Parity N (PRL=04, RLQ=01) Example

The allocation of data blocks in a Non-Rotating Parity N VD MUST adhere to the following formula:

Eq. 8

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), \text{MOD}(\text{FLOOR}(x/L), N-1), \text{FLOOR}(\text{FLOOR}(x/L)/N-1)).$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 9

$$\text{parity_block}(k, N-1, j) = \bigoplus_{i=0}^{N-2} \text{extent_block}(k, i, j).$$

All parity blocks MUST reside on extent $N-1$. Thus, i MUST equal $N-1$ for all parity blocks.

4.2.8 RAID-5 Rotating Parity 0 with Data Restart (PRL=05, RLQ=00)

Figure 9 gives an example of Rotating Parity 0 with Data Restart. Rotating Parity 0 with Data Restart is an implementation of RAID-5. Table 6 defines the indices and constants used in the description of RAID-5.

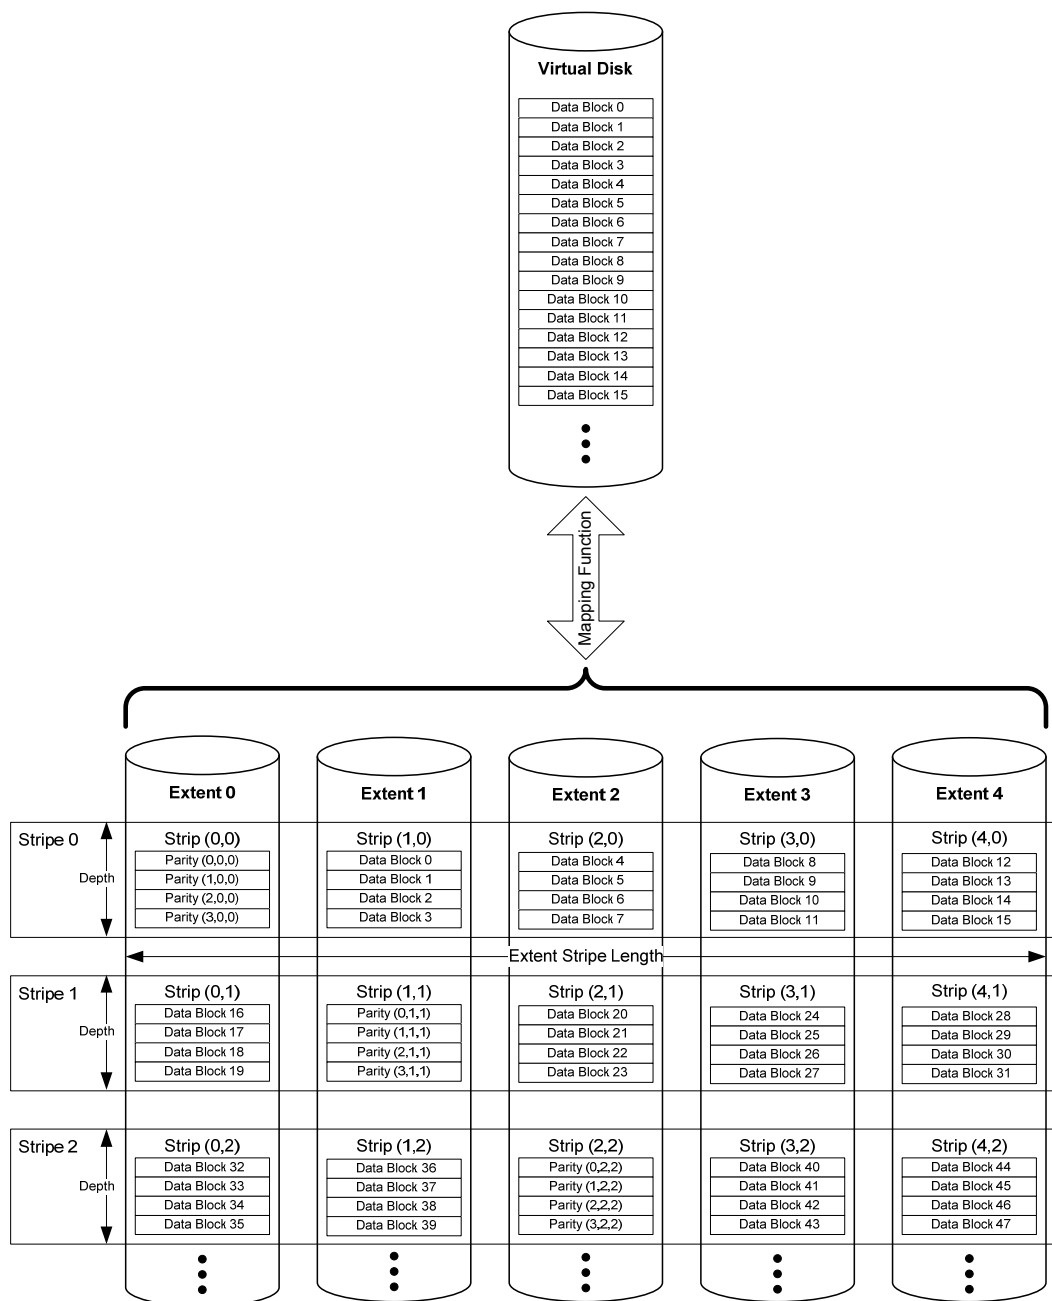


Figure 9: Rotating Parity 0 with Data Restart (PRL=05, RLQ=00) Example

Table 6: Indices and Constants for RAID-5 Rotating Parity 0 and N

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD	0	$N-1$
j	Index of a stripe in a VD	0	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/(N-1))$
k	Offset of a block from the beginning of a strip	0	$L-1$
L	Size of a strip in blocks	N/A (a fixed value)	N/A (a fixed value)
M	Number of data blocks in a VD. M MUST be evenly divisible by $(N-1)*L$	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD	N/A (a fixed value)	N/A (a fixed value)
x	Offset of a data block from the beginning of a VD	0	$M-1$
p	Index of the extent on which the parity_blocks for a given stripe reside	0	$N-1$

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 10

$$p = \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/L)/(N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 11

$$\begin{aligned} &\text{IF } \text{MOD}(\text{FLOOR}(x/L), N-1) < p \\ &\text{THEN } i = \text{MOD}(\text{FLOOR}(x/L), N-1) \\ &\text{ELSE } i = \text{MOD}(\text{FLOOR}(x/L), N-1) + 1. \end{aligned}$$

The allocation of data blocks in a Rotating Parity 0 with Data Restart VD MUST adhere to the following formula:

Eq. 12

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, \text{FLOOR}(\text{FLOOR}(x/L)/(N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 13

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

4.2.9 RAID-5 Rotating Parity N with Data Restart (PRL=05, RLQ=02)

Figure 10 gives an example of an implementation of RAID-5 called Rotating Parity N with Data Restart. The indices and constants listed in Table 6 are valid for this type of RAID.

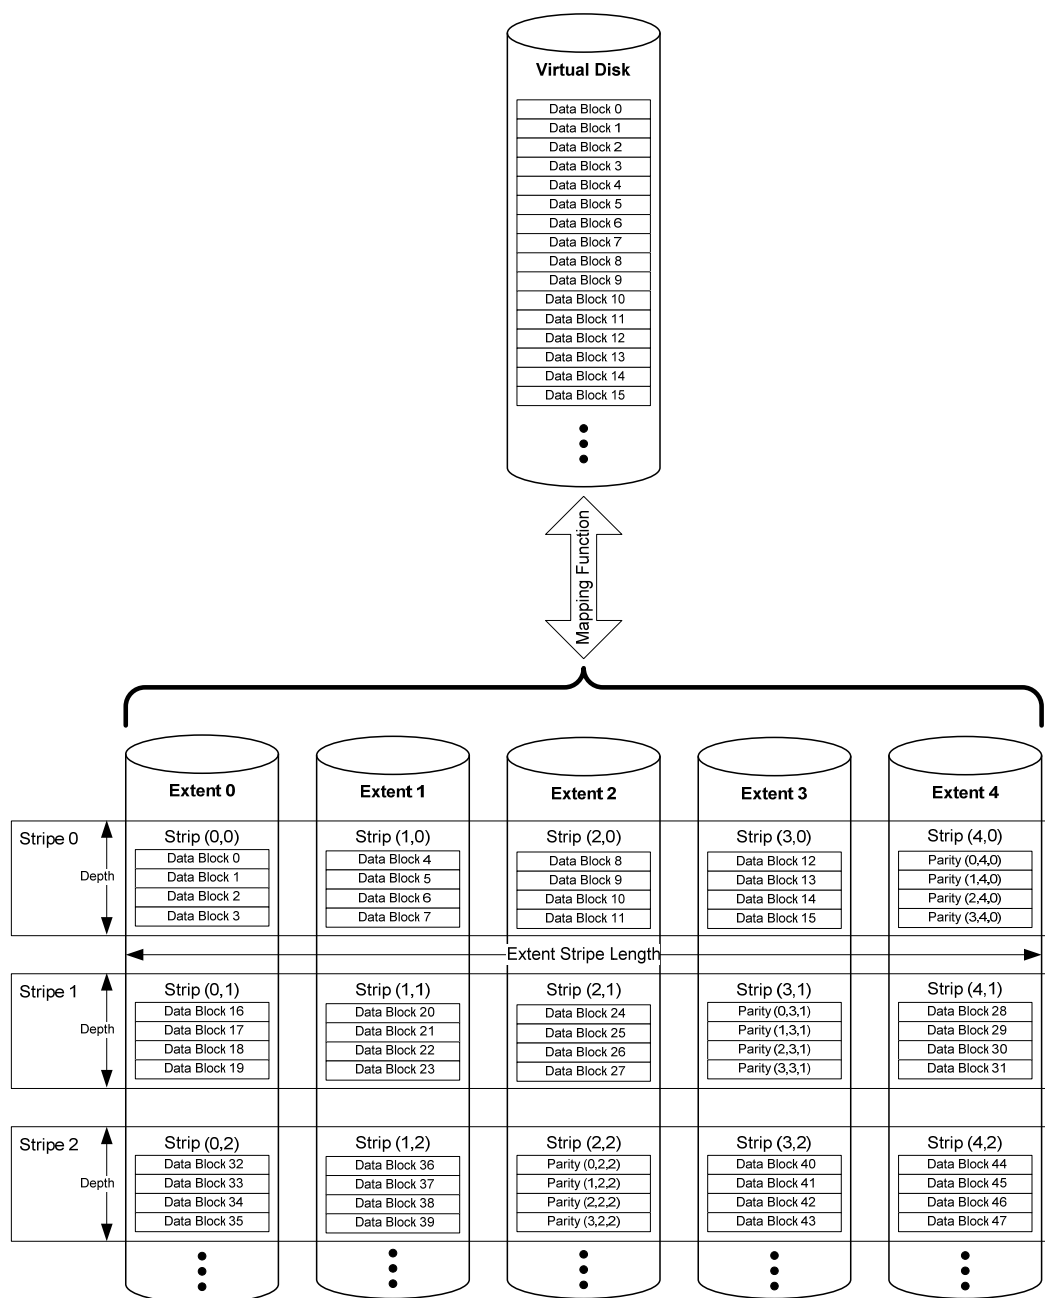


Figure 10: Rotating Parity N with Data Restart (PR=05, RLQ=02) Example

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 14

$$p = (N-1) - \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/L)/(N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 15

$$\begin{aligned} &\text{IF } \text{MOD}(\text{FLOOR}(x/L), N-1) < p \\ &\text{THEN } i = \text{MOD}(\text{FLOOR}(x/L), N-1) \\ &\text{ELSE } i = \text{MOD}(\text{FLOOR}(x/L), N-1) + 1. \end{aligned}$$

The allocation of data blocks in a Rotating Parity N with Data Restart VD MUST adhere to the following formula:

Eq. 16

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, \text{FLOOR}(\text{FLOOR}(x/L)/(N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 17

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

4.2.10 RAID-5 Rotating Parity N with Data Continuation (PRL=05, RLQ=03)

Figure 11 gives an example of RAID-5 implemented with Rotating Parity N with Data Continuation. The indices and constants given in Table 6 also apply to the formulas given below for Rotating Parity N with Data Continuation.

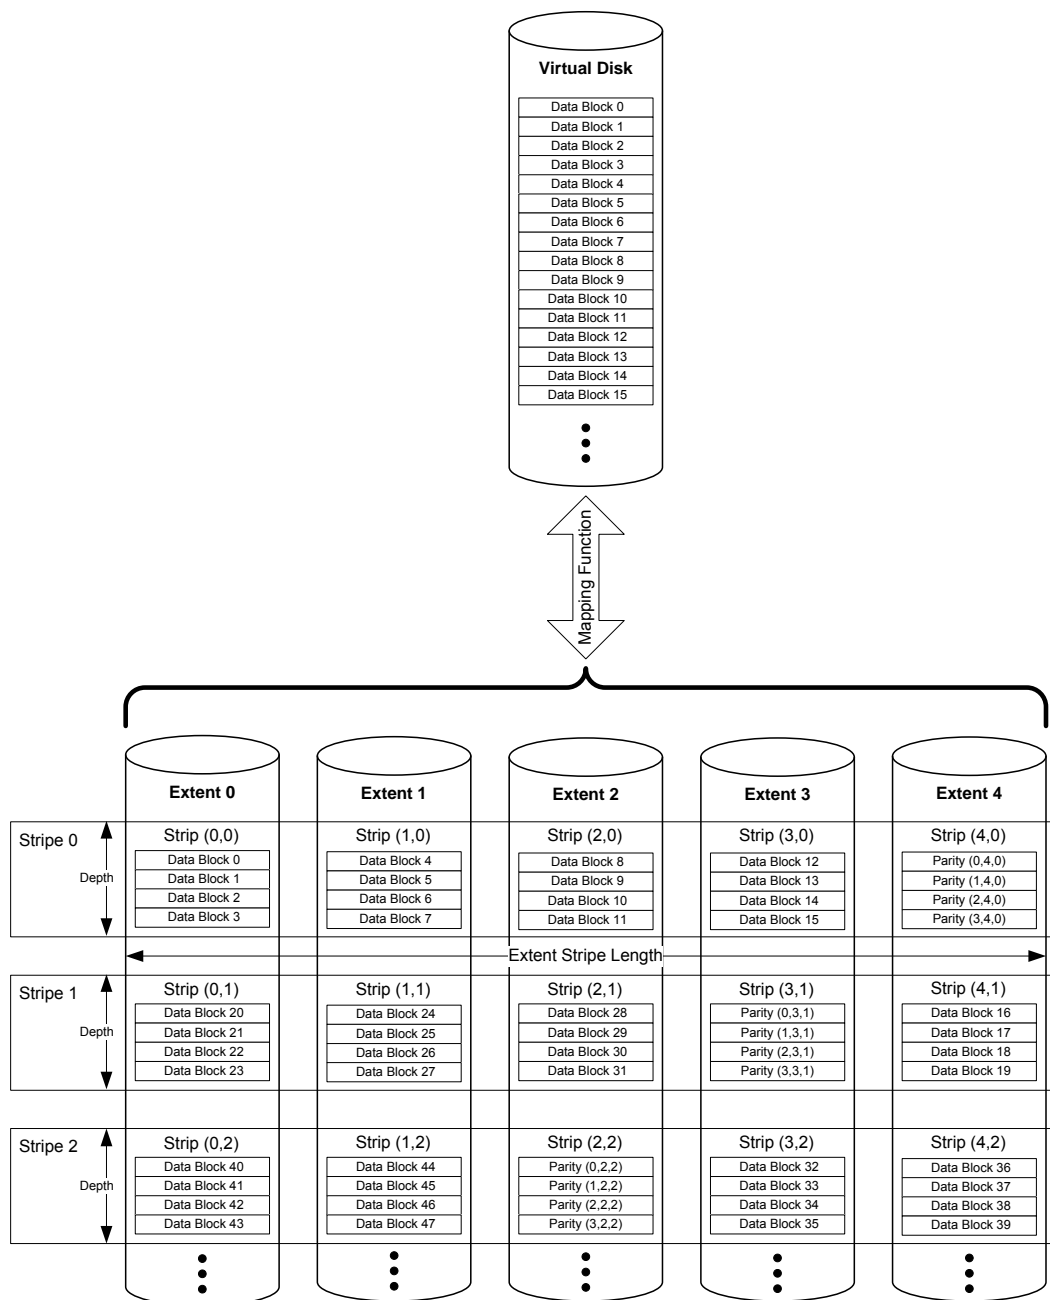


Figure 11: Rotating Parity N with Data Continuation (PRL=05, RLQ=03) Example

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 18

$$p = (N-1) - \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/L)/(N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 19

$$i = \text{MOD}(\text{MOD}(\text{FLOOR}(x/L), (N-1)) + p + 1, N).$$

The allocation of data blocks in a Rotating Parity N with Data Continuation VD MUST adhere to the following formula:

Eq. 20

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, \text{FLOOR}(\text{FLOOR}(x/L)/(N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 21

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

4.2.11 RAID-5E Rotating Parity 0 with Data Restart (PRL=15, RLQ=00)

Figure 12 gives an example of RAID-5E implemented with Rotating Parity 0 with Data Restart.

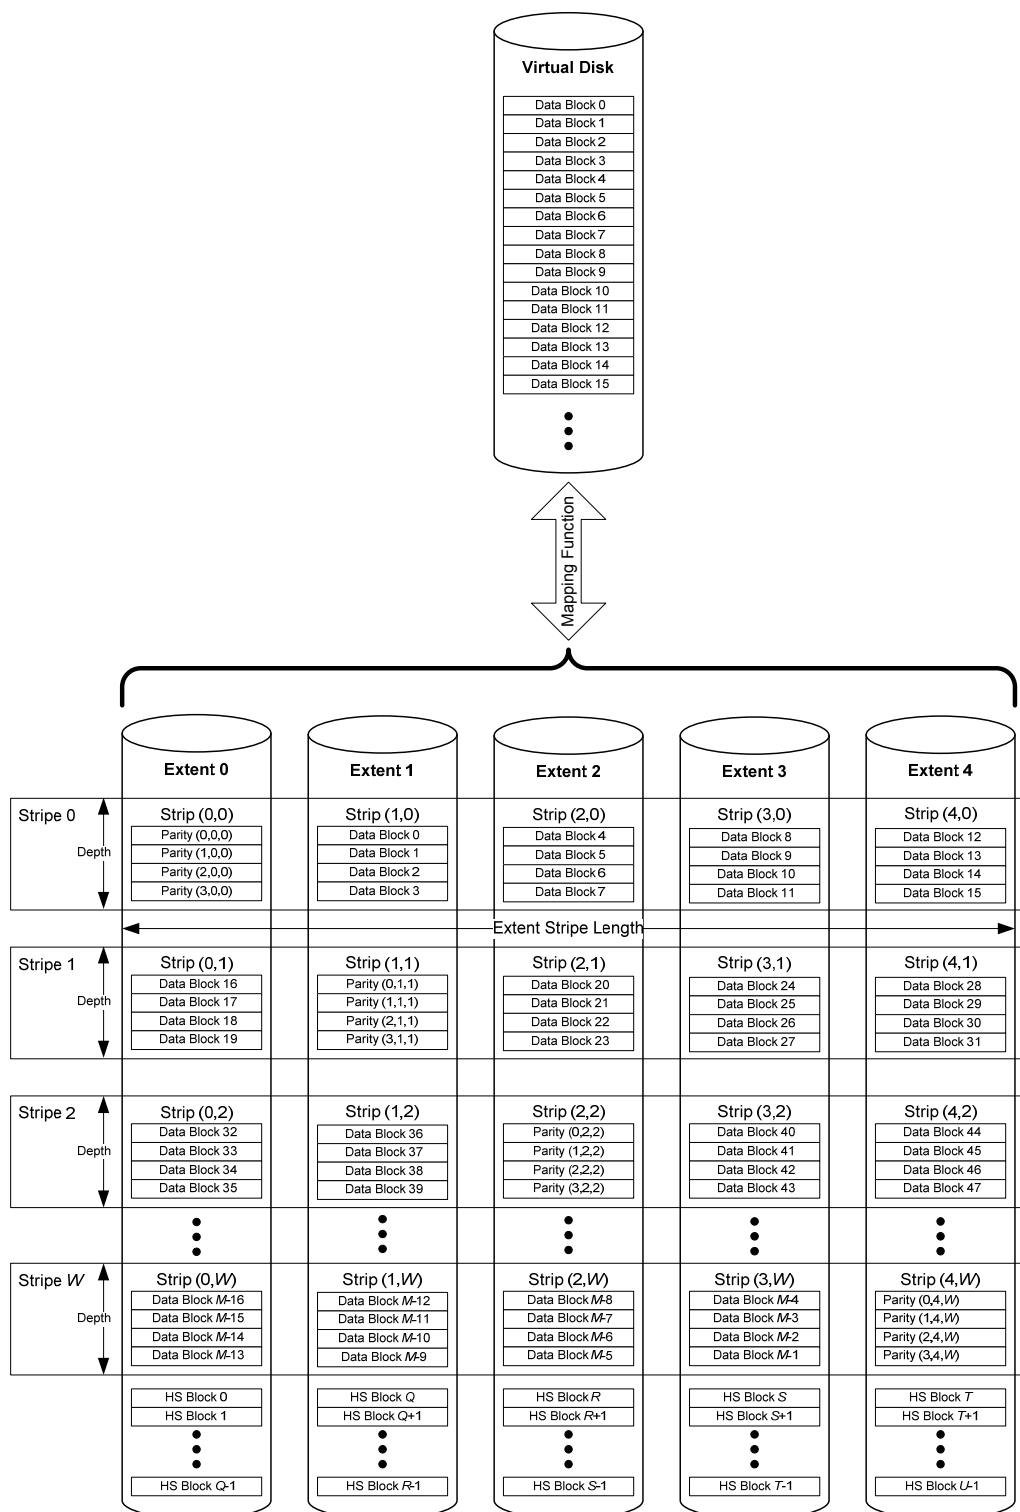


Figure 12: RAID-5E Rotating Parity 0 with Data Restart (PRL=15, RLQ=00) Example

RAID-5E has hot space at the end of each extent. In the event of an extent failure, the hot space on the remaining extents is used to rebuild and re-stripe the data in a manner that the remaining extents become a RAID-5 VD.

Table 7 gives the indices and constants used to describe the data layout of this type of RAID in the following formulas.

Table 7: Indices and Constants for RAID-5E

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD	0	$N-1$
j	Index of a stripe in a VD	0	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/(N-1))$
k	Offset of a block from the beginning of a strip	0	$L-1$
L	Size of a strip in blocks	N/A (a fixed value)	N/A (a fixed value)
M	Number of data blocks in a VD. M MUST be evenly divisible by $(N-1)*L$	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD	N/A (a fixed value)	N/A (a fixed value)
x	Offset of a data block from the beginning of a VD	0	$M-1$
p	Index of the extent on which the parity_blocks for a given stripe reside	0	$N-1$
U	Number of hot space blocks	$\geq ((M/(N-1)) * N) - M$	$\geq ((M/(N-1)) * N) - M$
W	Index of the last stripe containing data blocks	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/(N-1))$	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/(N-1))$

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 22

$$p = \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/L)/(N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 23

$$\begin{aligned} &\text{IF } \text{MOD}(\text{FLOOR}(x/L), N-1) < p \\ &\text{THEN } i = \text{MOD}(\text{FLOOR}(x/L), N-1) \\ &\text{ELSE } i = \text{MOD}(\text{FLOOR}(x/L), N-1) + 1. \end{aligned}$$

The allocation of data blocks in a RAID-5E Rotating Parity 0 with Data Restart VD MUST adhere to the following formula:

Eq. 24

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, \text{FLOOR}(\text{FLOOR}(x/L)/(N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 25

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

The number of hot space blocks U MUST adhere to the following formula:

Eq. 26

$$U \geq ((M/(N-1)) * N) - M$$

Hot space blocks MUST begin at offset M from the beginning of the VD. The total number of blocks in the VD MUST equal $M+U$. The hot space blocks MUST be evenly distributed across all extents. All hot space blocks on an extent MUST reside at the end of the extent.

In the event of an extent failure, the controller MUST reallocate the data as described in Section 4.2.8. The number of extents used in the resulting RAID-5 VD MUST be equal to the number of extents used in the RAID-5E VD reduced by one.

4.2.12 RAID-5E Rotating Parity N with Data Restart (PRL=15, RLQ=02)

Figure 13 gives an example of a RAID-5E implementation utilizing Rotating Parity N with Data Restart. The indices and constants listed in Table 7 are valid for this type of RAID.

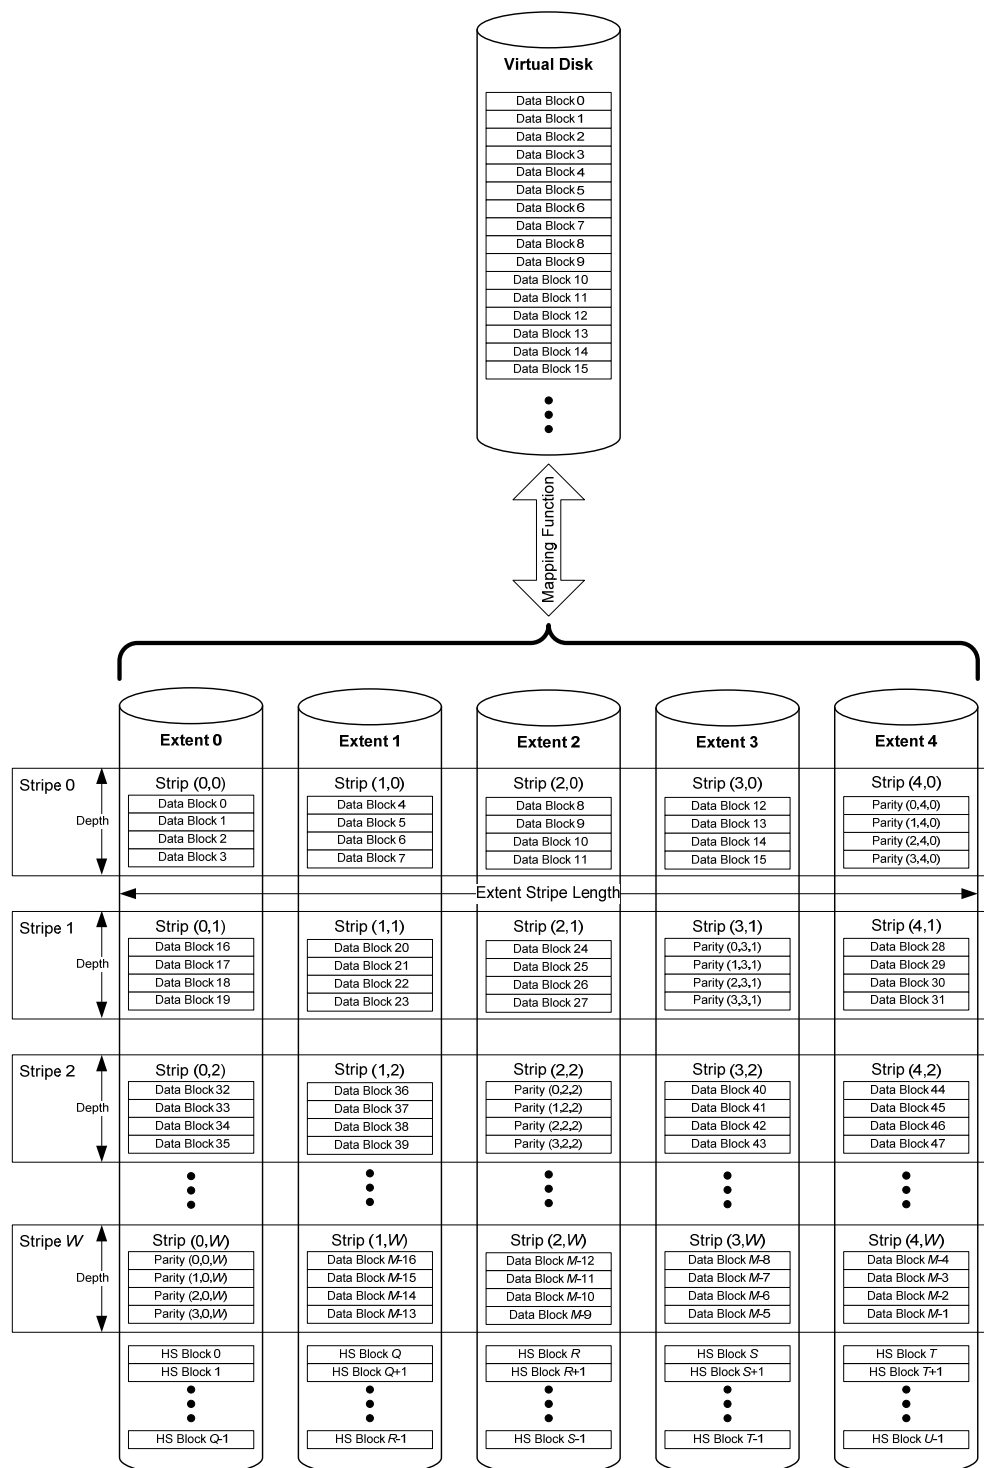


Figure 13: RAID-5E Rotating Parity N with Data Restart (PRL=15, RLQ=02) Example

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 27

$$p = (N-1) - \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/L)/(N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 28

$$\begin{aligned} &\text{IF } \text{MOD}(\text{FLOOR}(x/L), N-1) < p \\ &\text{THEN } i = \text{MOD}(\text{FLOOR}(x/L), N-1) \\ &\text{ELSE } i = \text{MOD}(\text{FLOOR}(x/L), N-1) + 1. \end{aligned}$$

The allocation of data blocks in a RAID-5E Rotating Parity N with Data Restart VD MUST adhere to the following formula:

Eq. 29

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, \text{FLOOR}(\text{FLOOR}(x/L)/(N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 30

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

The number of hot space blocks U MUST adhere to the following formula:

Eq. 31

$$U \geq ((M/(N-1)) * N) - M$$

Hot space blocks MUST begin at offset M from the beginning of the VD. The total number of blocks in the VD MUST equal $M+U$. The hot space blocks MUST be evenly distributed across all extents. All hot space blocks on an extent MUST reside at the end of the extent.

In the event of an extent failure, the controller MUST reallocate the data as described in Section 4.2.9. The number of extents used in the resulting RAID-5 VD MUST be equal to the number of extents used in the RAID-5E VD reduced by one.

4.2.13 RAID-5E Rotating Parity N with Data Continuation (PRL=15, RLQ=03)

Figure 14 gives an example of a RAID-5E volume implemented utilizing Rotating Parity N with Data Continuation. The indices and constants listed in Table 7 are valid for this type of RAID.

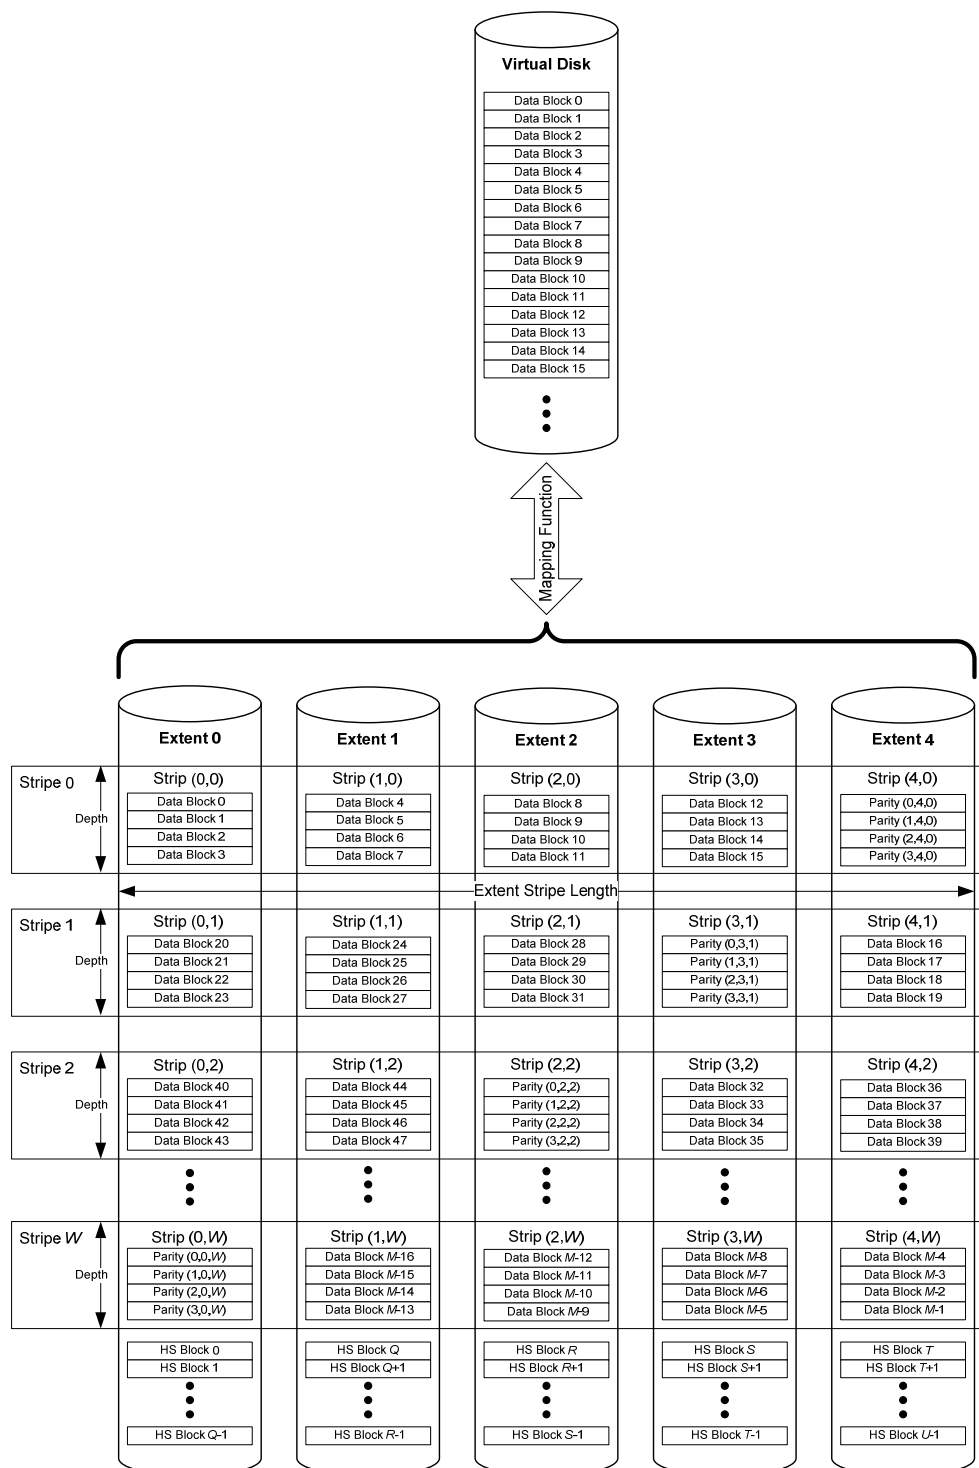


Figure 14: RAID-5E Rotating Parity N with Data Continuation (PRL=15, RLQ=03) Example

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 32

$$p = (N-1) - \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/L)/(N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 33

$$i = \text{MOD}(\text{MOD}(\text{FLOOR}(x/L), (N-1)) + p + 1, N).$$

The allocation of data blocks in a RAID-5E Rotating Parity N with Data Continuation VD MUST adhere to the following formula:

Eq. 34

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, \text{FLOOR}(\text{FLOOR}(x/L)/(N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 35

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

The number of hot space blocks U MUST adhere to the following formula:

Eq. 36

$$U \geq ((M/(N-1)) * N) - M$$

Hot space blocks MUST begin at offset M from the beginning of the VD. The total number of blocks in the VD MUST equal $M+U$. The hot space blocks MUST be evenly distributed across all extents. All hot space blocks on an extent MUST reside at the end of the extent.

In the event of an extent failure, the controller MUST reallocate the data as described in Section 4.2.10. The number of extents used in the resulting RAID-5 VD MUST be equal to the number of extents used in the RAID-5E VD reduced by one.

4.2.14 RAID-5EE Rotating Parity 0 with Data Restart (PRL=25, RLQ=00)

Figure 15 gives an example of RAID-5EE implemented with Rotating Parity 0 with Data Restart. RAID-5EE has hot space distributed across the extents. In the event of an extent failure, the hot space on the remaining extents is used to rebuild and re-stripe the data in a manner that the remaining extents become a RAID-5 VD. Table 8 gives the indices and constants used to describe the data layout of this type of RAID.

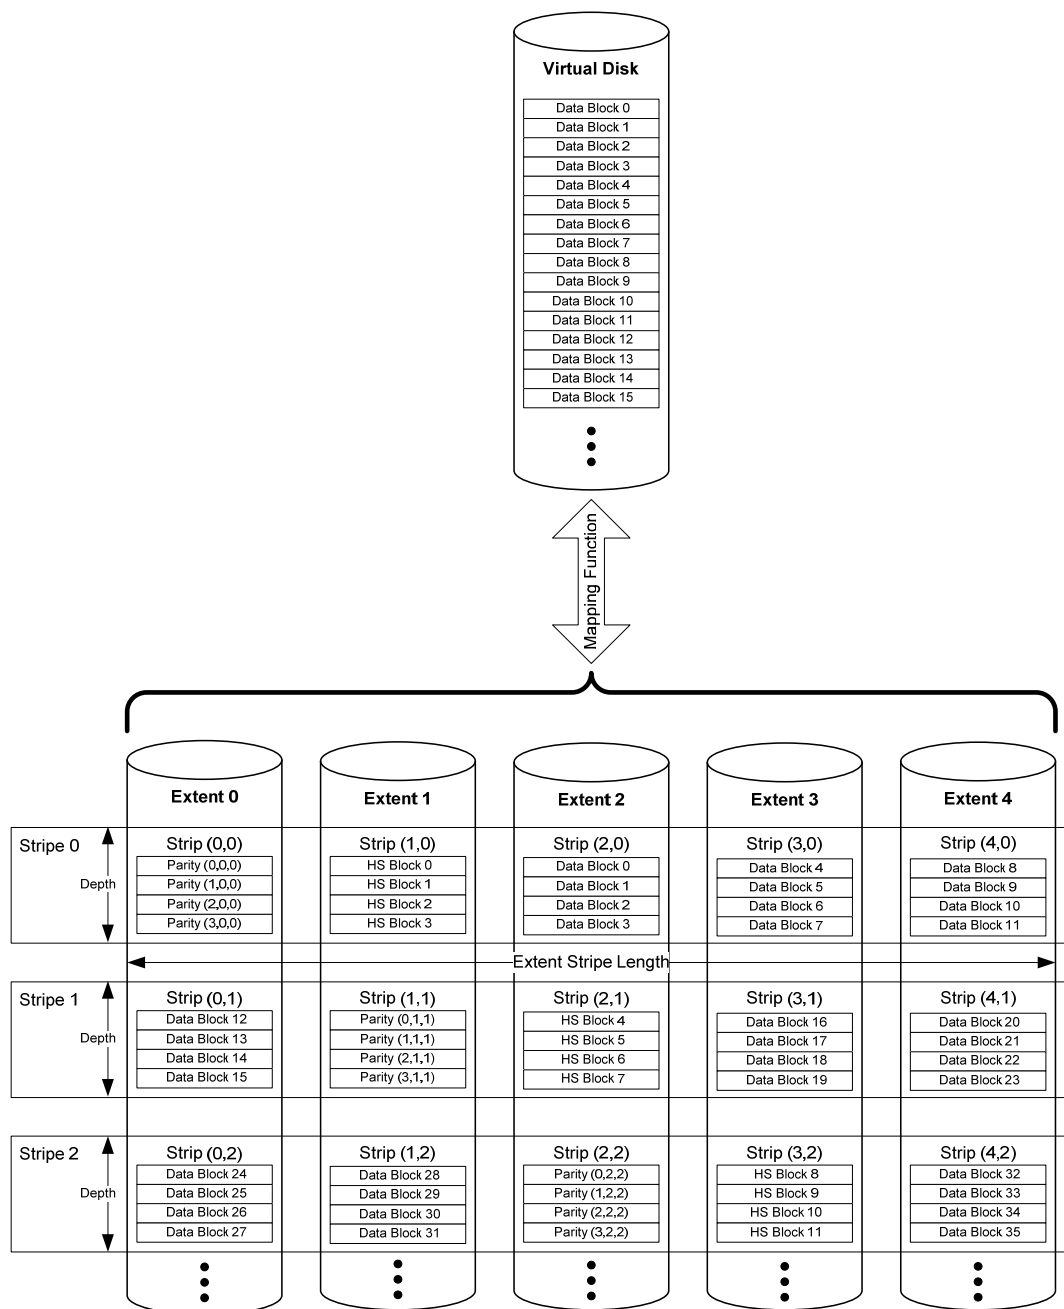


Figure 15: RAID-5EE Rotating Parity 0 with Data Restart (PRL=25, RLQ=00) Example

Table 8: Indices and Constants for RAID-5EE

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD	0	$N-1$
j	Index of a stripe in a VD	0	$\text{FLOOR}(\text{FLOOR}(M-1/L)/(N-2))$
k	Offset of a block from the beginning of a strip	0	$L-1$
L	Size of a strip in blocks	N/A (a fixed value)	N/A (a fixed value)
M	Number of data blocks in a VD. M MUST be evenly divisible by $(N-2)*L$	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD	N/A (a fixed value)	N/A (a fixed value)
x	Offset of a data block from the beginning of a VD	0	$M-1$
p	Index of the extent on which the parity_blocks for a given stripe reside	0	$N-1$
h	Index of the extent on which the hot space for a given stripe resides	0	$N-1$

The stripe j on which a given virtual block x resides MUST adhere to the following formula:

Eq. 37

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-2))$$

The extent p on which the parity for a given stripe j resides MUST adhere to the following formula:

Eq. 38

$$p = \text{MOD}(j, N).$$

The extent h on which hot space for a given stripe j resides MUST adhere to the following formula:

Eq. 39

$$h = \text{MOD}((p+1), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 40

IF $\text{MOD}(\text{FLOOR}(x/L), N-2) < p$

THEN

IF $\text{MOD}(\text{FLOOR}(x/L), N-2) < h$

THEN

$$i = \text{MOD}(\text{FLOOR}(x/L), N-2)$$

ELSE

$$i = \text{MOD}(\text{FLOOR}(x/L), N-2) + 1$$

ELSE

$$i = \text{MOD}(\text{FLOOR}(x/L, 1), N-2) + 2$$

The allocation of data blocks in a RAID-5EE Rotating Parity 0 with Data Restart VD MUST adhere to the following formula:

Eq. 41

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, j)$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 42

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p, i \neq h} \text{extent_block}(k, i, j)$$

In the event of an extent failure, the controller MUST reallocate the data as describe in Section 4.2.8. The number of extents used in the resulting RAID-5 VD MUST be equal to the number of extents used in the RAID-5EE volume reduced by one.

4.2.15 RAID-5EE Rotating Parity N with Data Restart (PRL=25, RLQ=02)

Figure 16 gives an example of RAID-5EE implemented with Rotating Parity N with Data Restart. The indices and constants given in Table 8 are used in the description below.

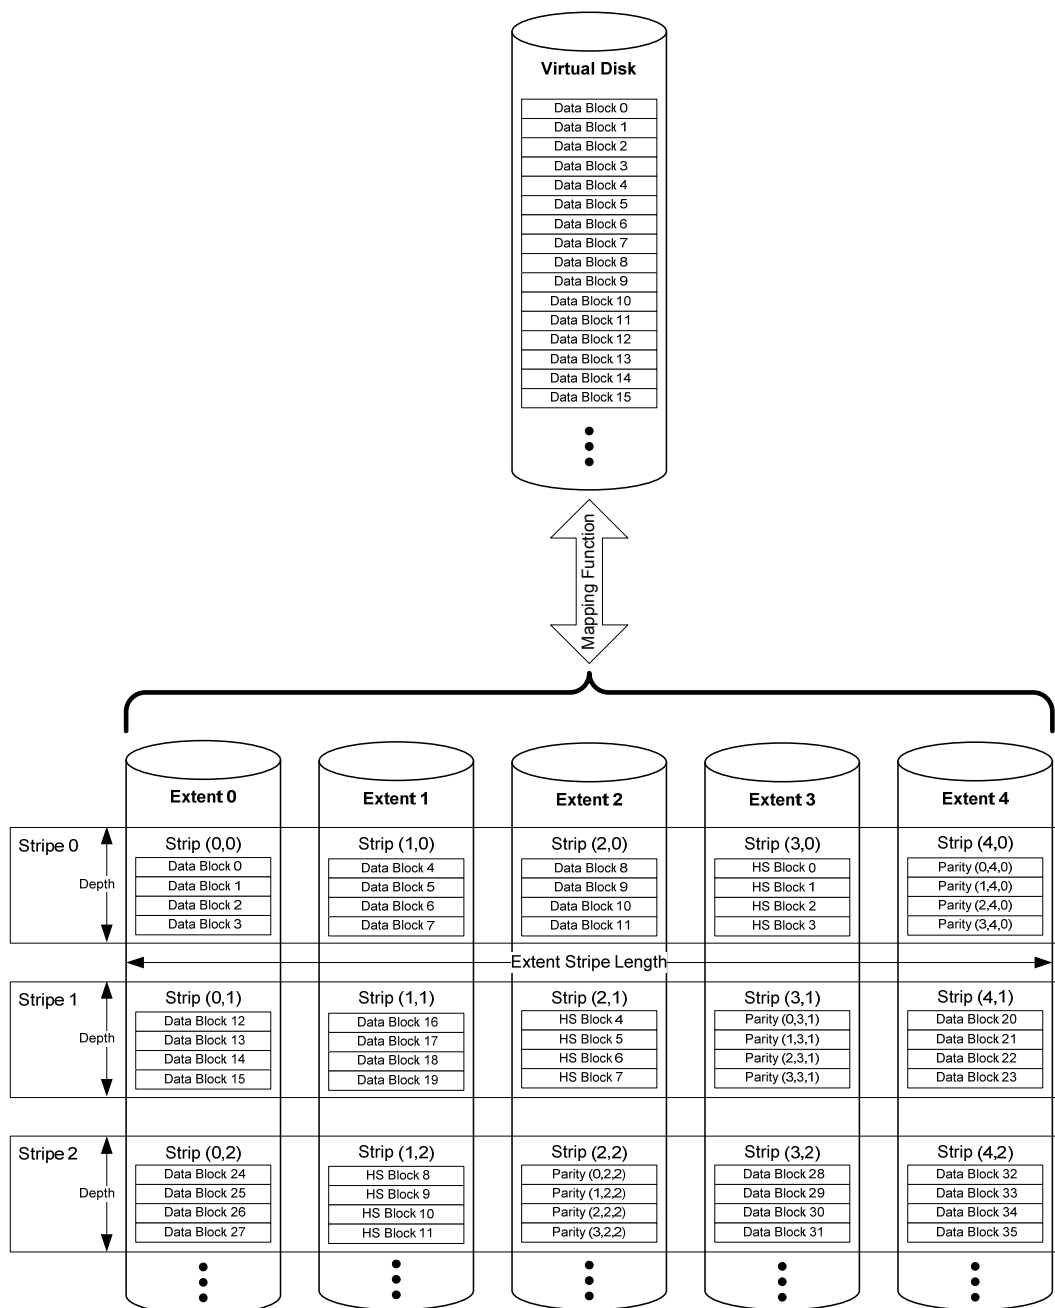


Figure 16: RAID-5EE Rotating Parity N with Data Restart (PRL=25, RLQ=02) Example

The stripe j on which a given virtual block x resides MUST adhere to the following formula:

Eq. 43

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-2)).$$

The extent p on which the parity for a given stripe j resides MUST adhere to the following formula:

Eq. 44

$$p = (N-1) - \text{MOD}(j, N).$$

The extent h on which hot space for a given stripe j resides MUST adhere to the following formula:

Eq. 45

$$h = \text{MOD}((p-1), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 46

```

IF MOD(FLOOR(x/L), N-2) < h
THEN
    IF MOD(FLOOR(x/L), N-2) < p
    THEN
        i = MOD(FLOOR(x/L), N-2)
    ELSE
        i = MOD(FLOOR(x/L), N-2)+1
    ELSE
        i = MOD(FLOOR(x/L, 1), N-2)+2

```

The allocation of data blocks in a RAID-5EE Rotating Parity N with Data Restart VD MUST adhere to the following formula:

Eq. 47

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, j).$$

The values of the party blocks MUST be calculated according to the following formula:

Eq. 48

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p, i \neq h} \text{extent_block}(k, i, j).$$

In the event of an extent failure, the controller MUST reallocate the data as describe in Section 4.2.9. The number of extents used in the resulting RAID-5 VD MUST be equal to the number of extents used in the RAID-5EE volume reduced by one.

4.2.16 RAID-5EE Rotating Parity N with Data Continuation (PRL=25, RLQ=03)

Figure 17 gives an example of RAID-5EE implemented with Rotating Parity N with Data Continuation. The indices and constants given in Table 8 are used in the description below.

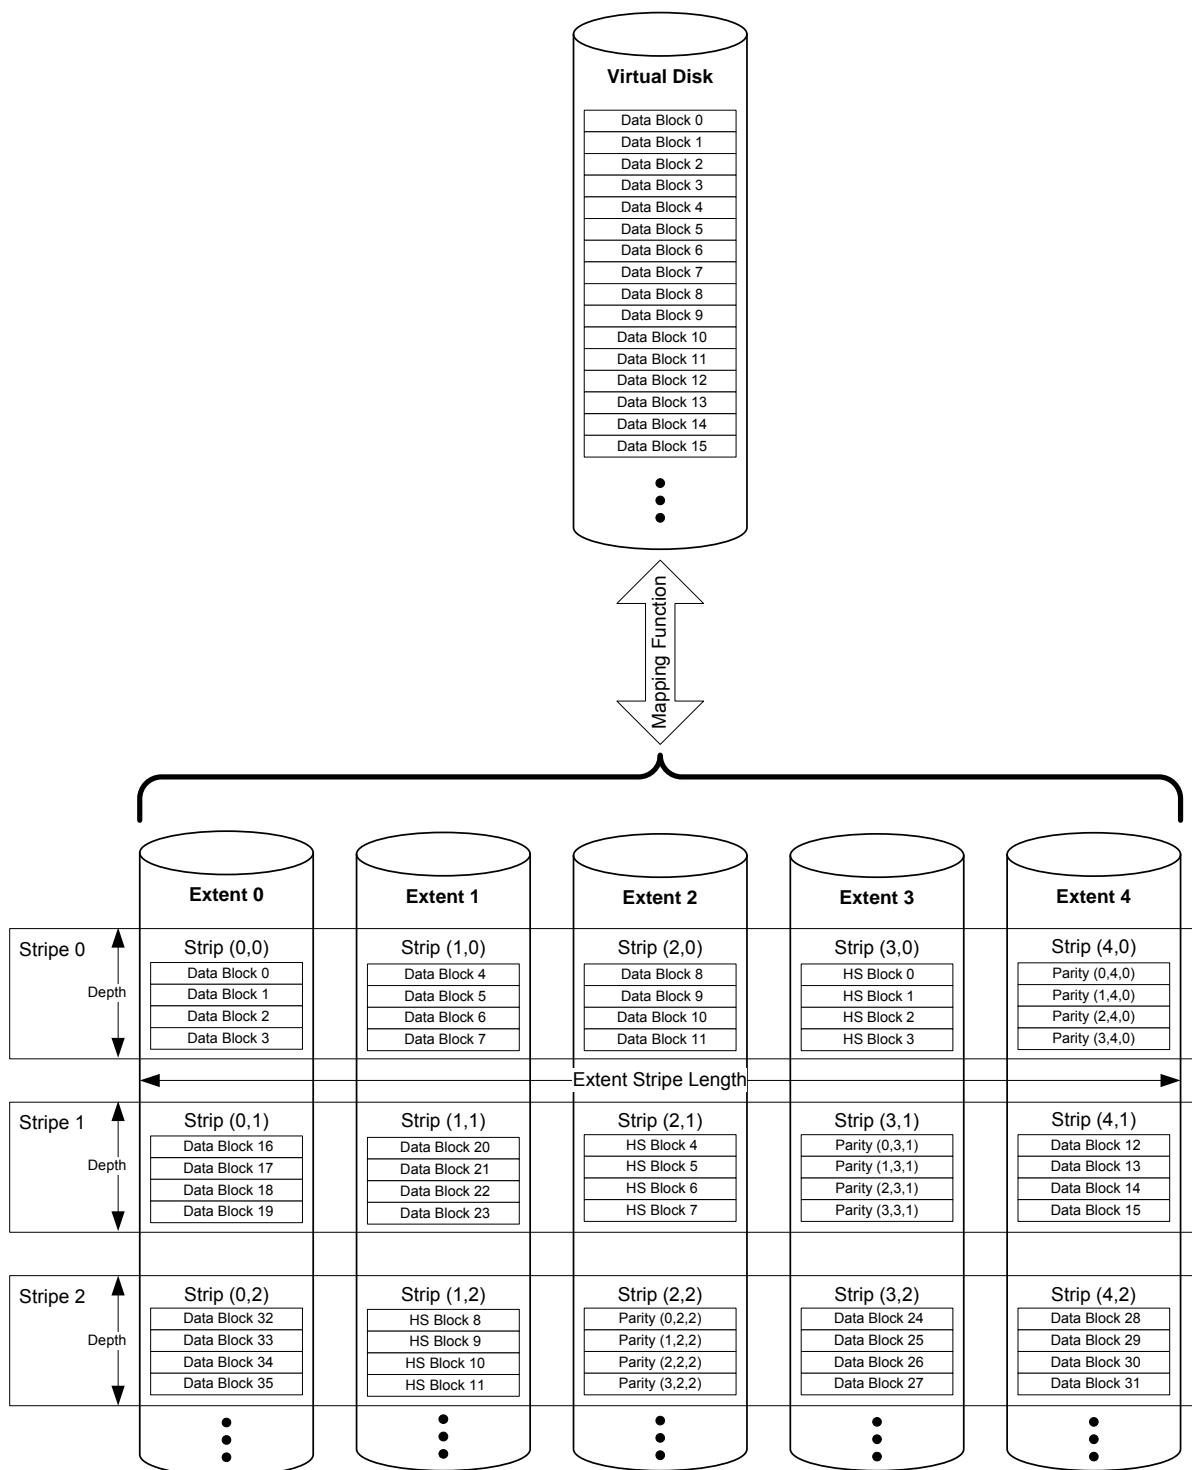


Figure 17: RAID-5EE Rotating Parity N with Data Continuation (PRL=25, RLQ=03) Example

The strip j on which a given virtual block x resides MUST adhere to the following formula:

Eq. 49

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-2)).$$

The extent p on which the parity for a given stripe j resides MUST adhere to the following formula:

Eq. 50

$$p = (N-1) - \text{MOD}(j, N).$$

The extent h on which hot space for a given stripe j resides MUST adhere to the following formula:

Eq. 51

$$h = \text{MOD}((p-1), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 52

$$i = \text{MOD}(\text{MOD}(\text{FLOOR}(x/L), N-2) + p + 1, N).$$

The allocation of data blocks in a RAID-5EE Rotating Parity N with Data Restart VD MUST adhere to the following formula:

Eq. 53

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), i, j).$$

The values of the party blocks MUST be calculated according to the following formula:

Eq. 54

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1-i-p-h} \text{extent_block}(k, i, j).$$

In the event of an extent failure, the controller MUST reallocate the data as describe in Section 4.2.10. The number of extents used in the resulting RAID-5 VD MUST be equal to the number of extents used in the RAID-5EE volume reduced by one.

4.2.17 RAID-5R Rotating Parity 0 after R Stripes with Data Restart (PRL=35, RLQ=00)

Figure 18 gives an example of RAID-5R Rotating Parity 0 after R Stripes with Data Restart. Table 9 gives the indices and constants used to describe the data layout of this type of RAID in the following formulas.

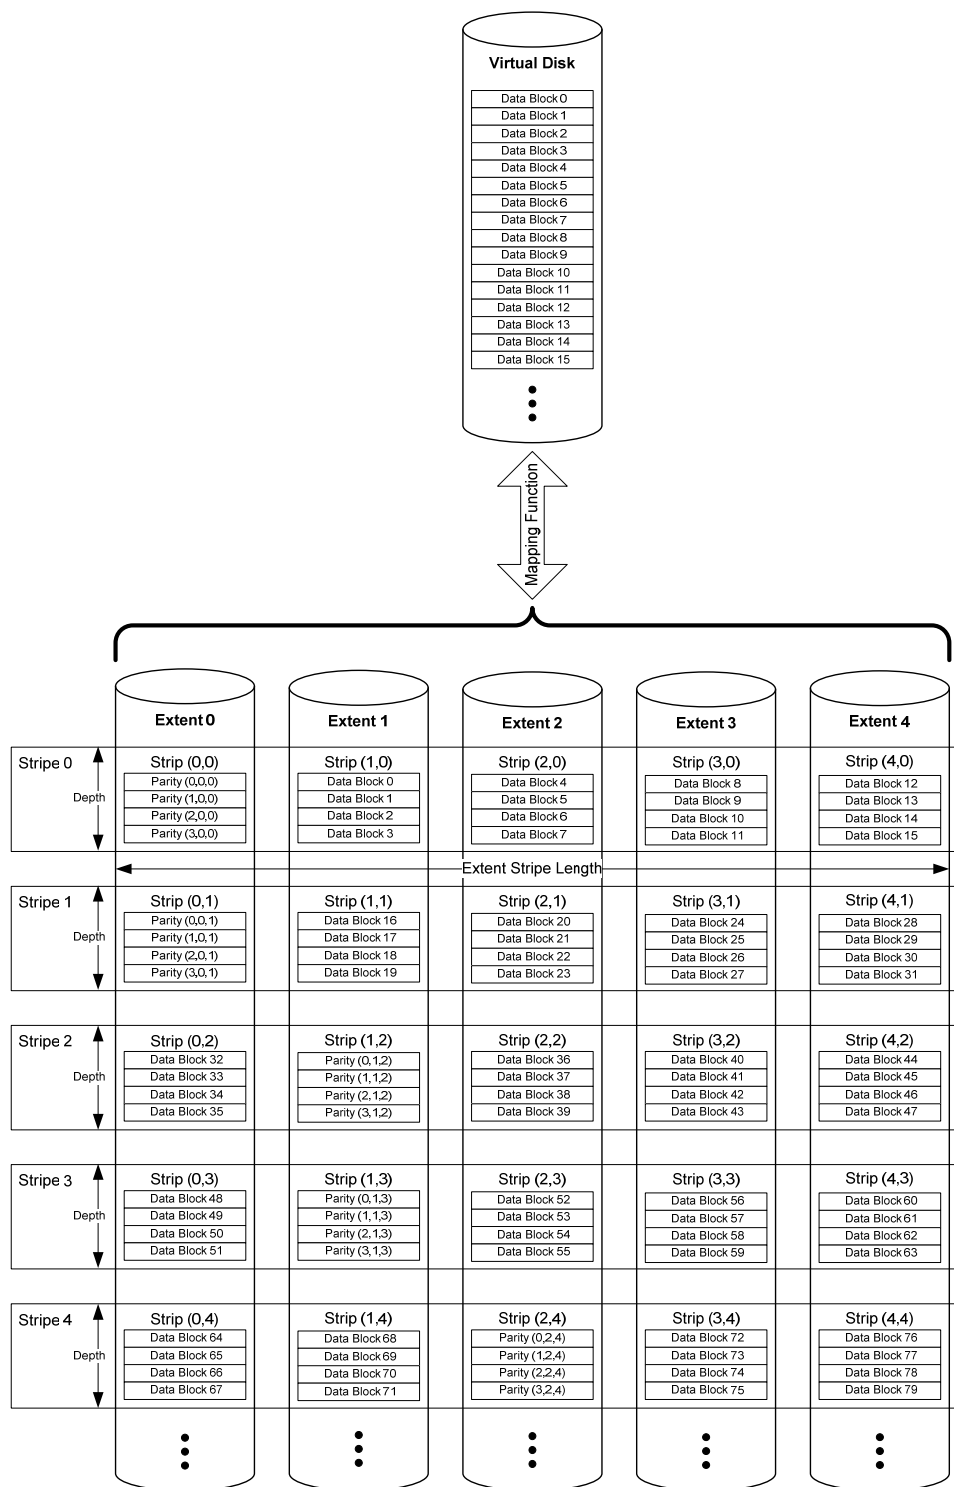


Figure 18: RAID-5R Rotating Parity 0 after R Stripes with Data Restart (PRL=35, RLQ=00) Example

Table 9: Indices and Constants for RAID-5R Rotating Parity 0 and N

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD	0	$N-1$
j	Index of a stripe in a VD	0	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/(N-1))$
k	Offset of a block from the beginning of a strip	0	$L-1$
L	Size of a strip in blocks	N/A (a fixed value)	N/A (a fixed value)
M	Number of data blocks in a VD. M MUST be evenly divisible by $(N-1)*L*R$	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD	N/A (a fixed value)	N/A (a fixed value)
R	Number of stripes before rotating Parity (where $R=2^n$)	1 (default)	0x80000000
x	Offset of a data block from the beginning of a VD	0	$M-1$
p	Index of the extent on which the parity_blocks for a given stripe reside	0	$N-1$

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 55

$$p = \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/(L*R))/(N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 56

$$\begin{aligned} &\text{IF } \text{MOD}(\text{FLOOR}(x/(L*R)), N-1) < p \\ &\quad \text{THEN } i = \text{MOD}(\text{FLOOR}(x/(L*R)), N-1) \\ &\quad \text{ELSE } i = \text{MOD}(\text{FLOOR}(x/(L*R)), N-1) + 1. \end{aligned}$$

The allocation of data blocks in a Rotating Parity 0 with Data Restart VD MUST adhere to the following formula:

Eq. 57

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), i, \text{FLOOR}(\text{FLOOR}(x/L)/(N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 58

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

4.2.18 RAID-5R Rotating Parity N after R Stripes with Data Restart (PRL=35, RLQ=02)

Figure 19 gives an example of an implementation of RAID-5 Rotating Parity N after R Stripes with Data Restart. The indices and constants listed in Table 9 are valid for this type of RAID.

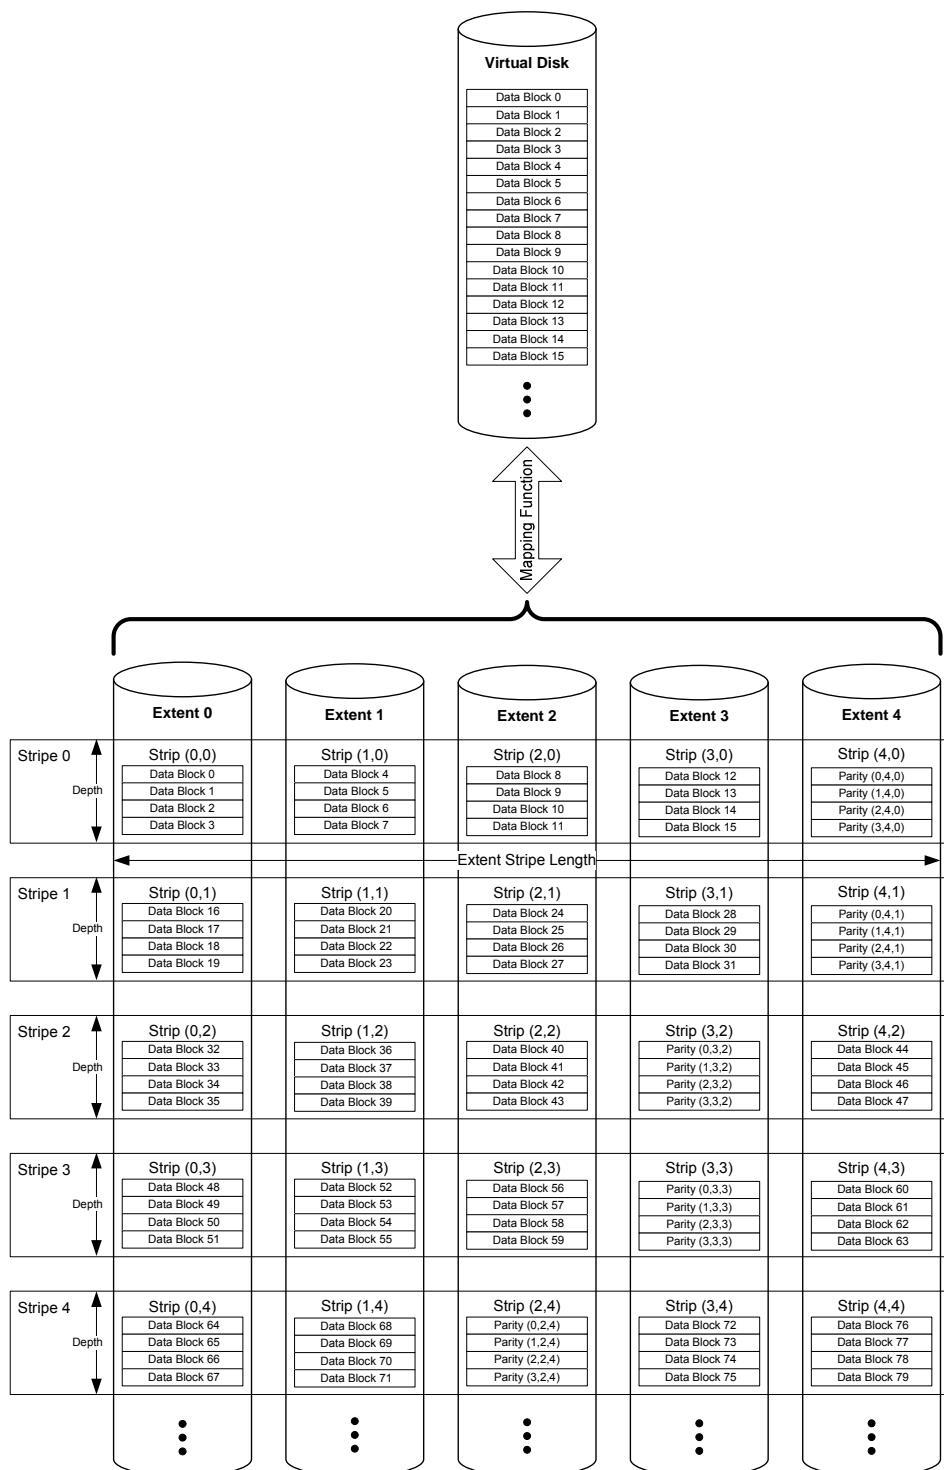


Figure 19: RAID-5R Rotating Parity N after R Stripes with Data Restart (PR=35, RLQ=02) Example

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 59

$$p = (N-1) - \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/(L * R)) / (N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 60

$$\begin{aligned} &\text{IF } \text{MOD}(\text{FLOOR}(x/(L * R)), N-1) < p \\ &\quad \text{THEN } i = \text{MOD}(\text{FLOOR}(x/(L * R)), N-1) \\ &\quad \text{ELSE } i = \text{MOD}(\text{FLOOR}(x/(L * R)), N-1) + 1. \end{aligned}$$

The allocation of data blocks in a Rotating Parity N with Data Restart VD MUST adhere to the following formula:

Eq. 61

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), i, \text{FLOOR}(\text{FLOOR}(x/L) / (N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 62

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

4.2.19 RAID-5R Rotating Parity N after R Stripes with Data Continuation (PRL=35, RLQ=03)

Figure 20 gives an example of RAID-5 Rotating Parity N after R Stripes with Data Continuation. The indices and constants given in Table 9 also apply to the formulas given below.

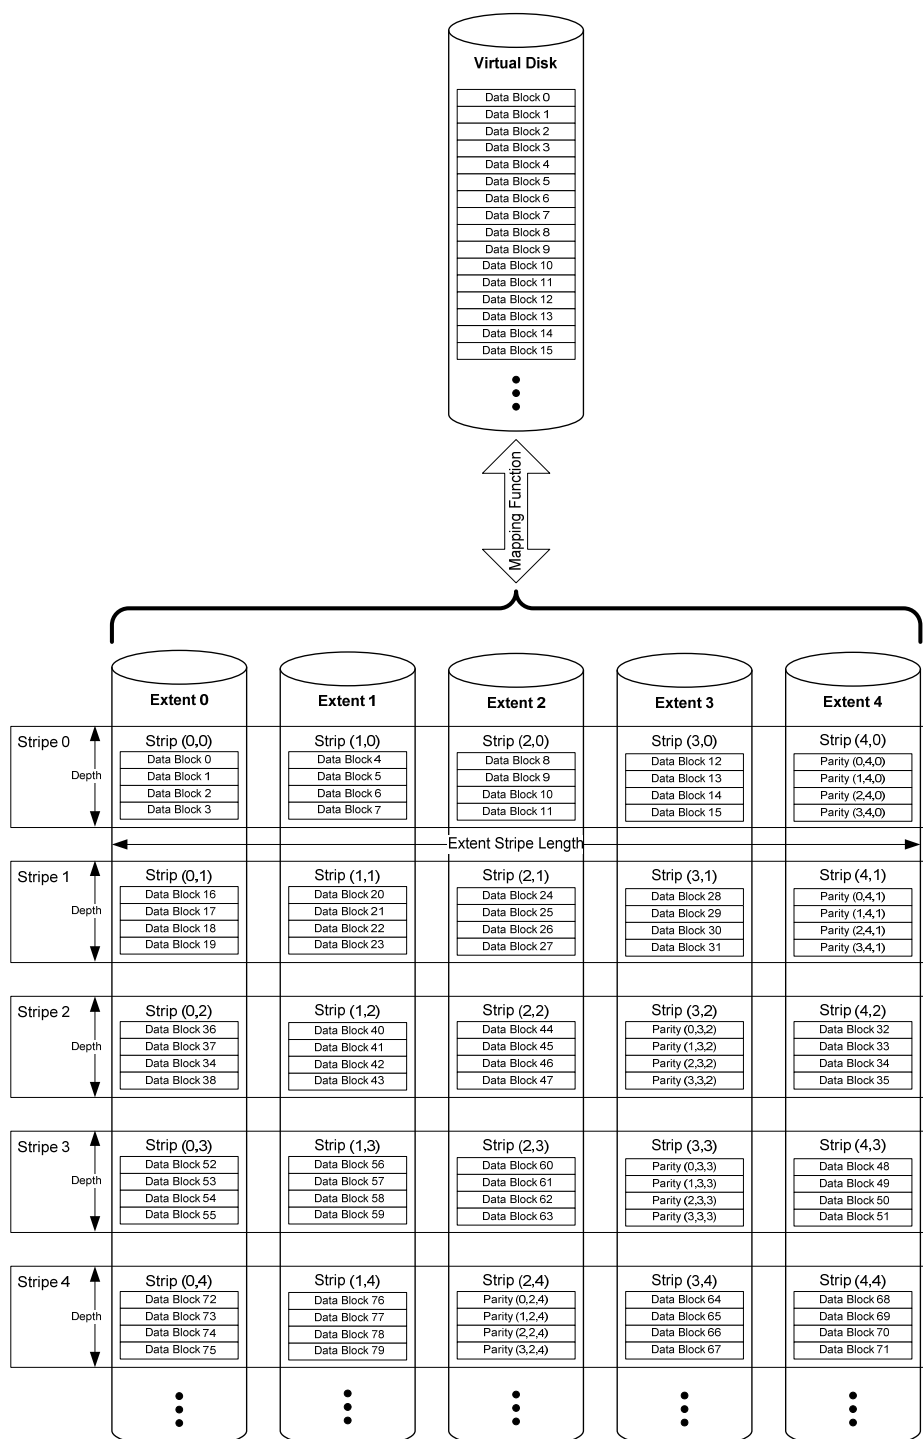


Figure 20: RAID-5R Rotating Parity N after R Stripes with Data Continuation (PRL=35, RLQ=03) Example

The extent p on which the parity block for a given virtual block x resides MUST adhere to the following formula:

Eq. 63

$$p = (N-1) - \text{MOD}(\text{FLOOR}(\text{FLOOR}(x/(L * R))/(N-1)), N).$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 64

$$i = \text{MOD}(\text{MOD}(\text{FLOOR}(x/(L * R)), (N-1)) + p + 1, N).$$

The allocation of data blocks in a Rotating Parity N with Data Continuation VD MUST adhere to the following formula:

Eq. 65

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), i, \text{FLOOR}(\text{FLOOR}(x/L)/(N-1)))$$

The values of the parity blocks MUST be calculated according to the following formula:

Eq. 66

$$\text{parity_block}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p} \text{extent_block}(k, i, j).$$

4.2.20 RAID-1E Integrated Adjacent Stripe Mirroring (PRL= 11, RLQ=00)

Figure 21 gives an example of Integrated Adjacent Strip Mirroring. Table 10 gives the indices and constants used to describe the data layout of this type of RAID in the following formulas.

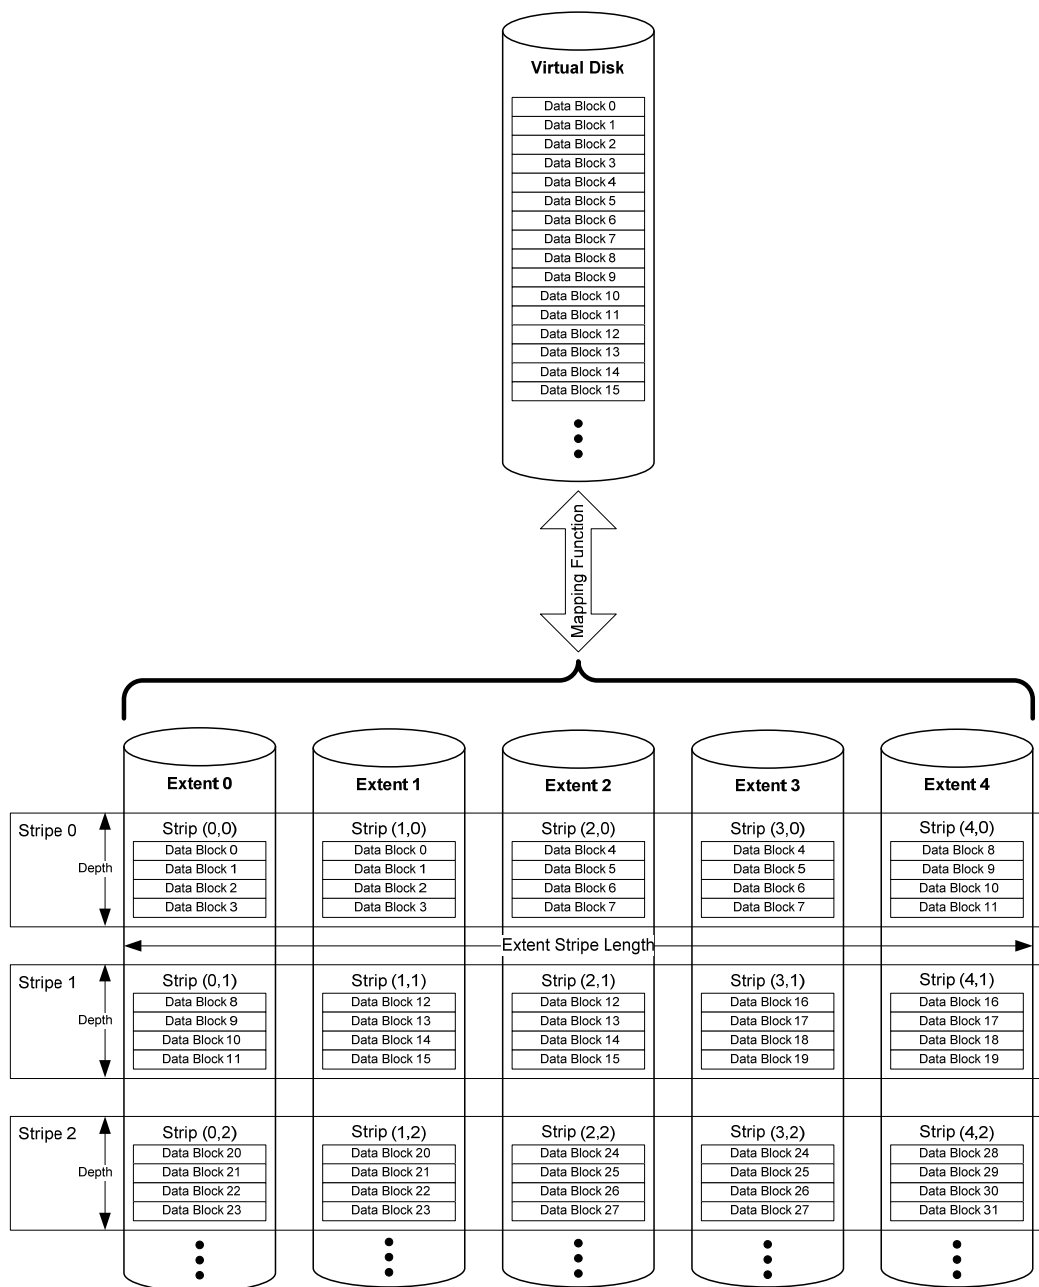


Figure 21: RAID-1E Integrated Adjacent Stripe Mirroring (PRL= 11, RLQ=00) Example

Table 10: Indices and Constants for Integrated Adjacent Strip Mirroring

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD where the primary copy of a virtual_block is stored.	0	$N-1$
j	Index of a stripe in a VD	0	$(M/N)-1$
k	Offset of a block from the beginning of a strip	0	$L-1$
L	Size of a strip in blocks	N/A (a fixed value)	N/A (a fixed value)
M	Number of data blocks in a VD. M MUST be evenly divisible by $N-1$.	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD	N/A (a fixed value)	N/A (a fixed value)
q	Index of an extent of a VD where the mirrored copy of a virtual_block is stored.	0	$N-1$
r	Index of the stripe where the mirrored copy of a virtual_block is stored	0	$(M/N)-1$
x	Offset of a data block from the beginning of a VD	0	$M-1$

In an Integrated Adjacent Stripe Mirroring VD, each virtual_block(x) MUST be stored in two locations. These two locations are referred to as extent_block (k, i, j) and extent_block' (k, q, r). The allocation of data in an Integrated Adjacent Stripe Mirroring VD MUST adhere to the following formulas:

Eq. 67

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), \text{MOD}(\text{FLOOR}(x/L)*2, N), \text{FLOOR}(\text{FLOOR}(x/L)*2/N))$$

Eq. 68

$$\text{virtual_block}(x) = \text{extent_block}'(\text{MOD}(x/L), \text{MOD}((\text{FLOOR}(x/L)*2)+1, N), \text{FLOOR}(((\text{FLOOR}(x/L)*2)+1)/N))$$

4.2.21 RAID-1E Integrated Offset Stripe Mirroring (PRL=11, RLQ=01)

Figure 22 gives an example of Integrated Offset Stripe Mirroring. The constants and indices given in Table 10 also apply for the following formulas that describe the data layout of an Integrated Offset Stripe Mirroring VD.

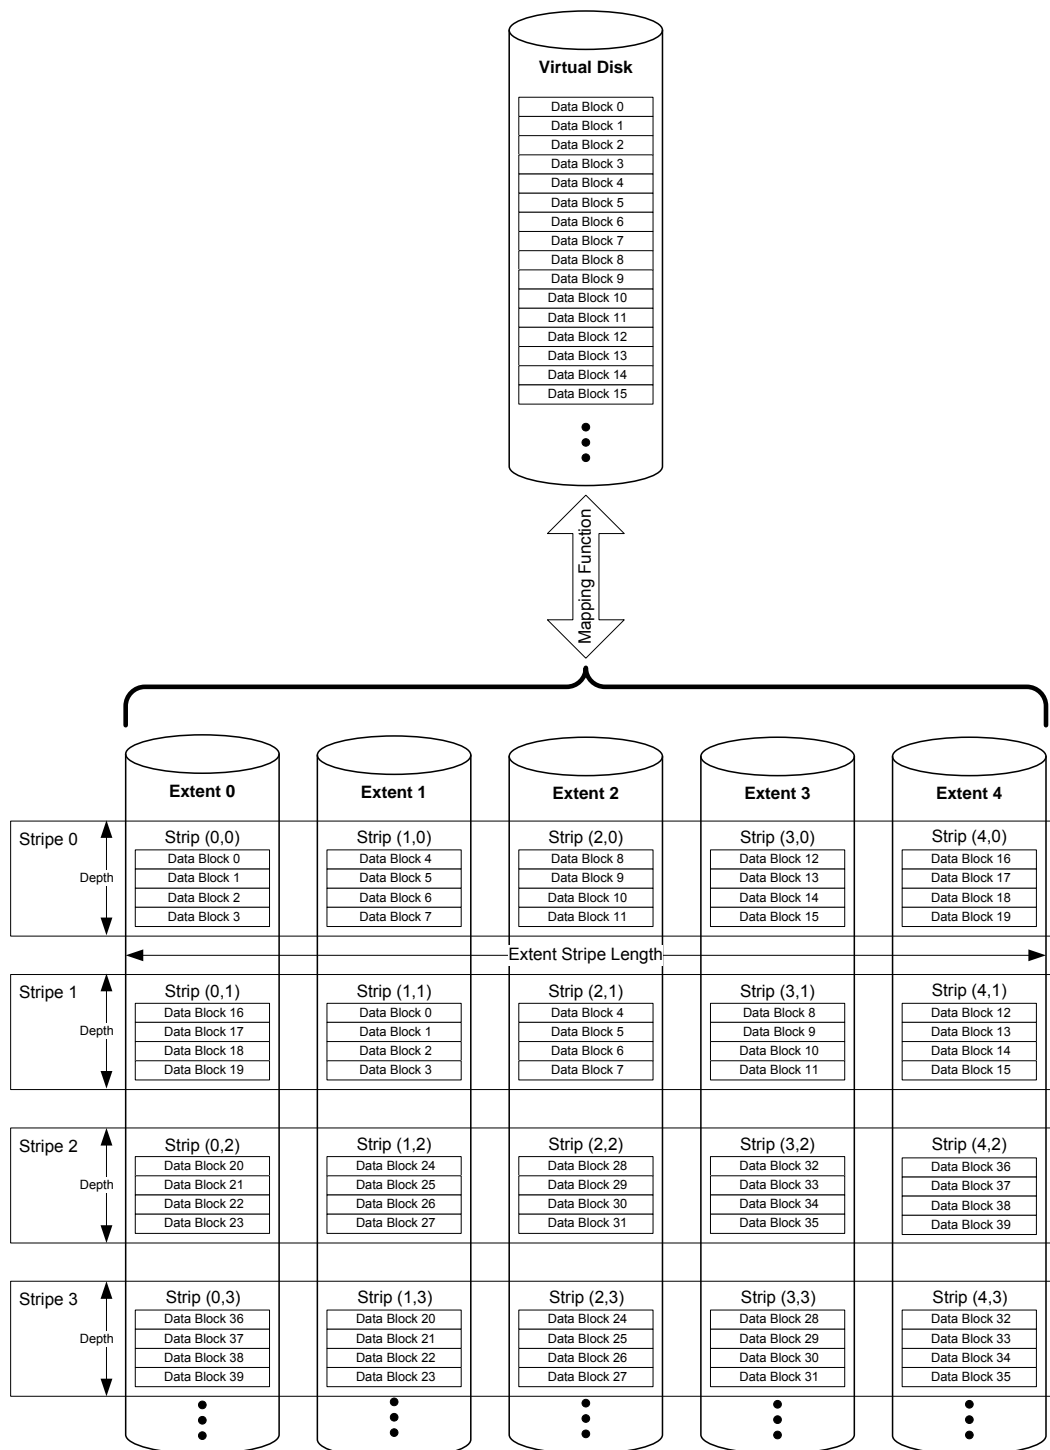


Figure 22: RAID-1E Integrated Offset Stripe Mirroring (PRL= 11, RLQ=01) Example

Each virtual_block (x) MUST be stored in two locations. These two locations are referred to as extent_block (k, i, j) and extent_block' (k, q, r). The allocation of data in an Integrated Offset Stripe Mirroring VD MUST adhere to the following formulas:

Eq. 69

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x/L), \text{MOD}(\text{FLOOR}(x/L), N), \text{FLOOR}(\text{FLOOR}(x/L)/N)*2)$$

Eq. 70

$$\text{virtual_block}(x) = \text{extent_block}'(\text{MOD}(x/L), \text{MOD}(\text{FLOOR}(x/L)+1, N), (\text{FLOOR}(\text{FLOOR}(x/L)/N)*2)+1)$$

4.2.22 RAID-6 Rotating Parity 0 with Data Restart (PRL=06, RLQ=01)

Figure 23 gives an example of RAID-6 called Rotating Parity 0 with Data Restart. Table 11 defines the indices and constants that are used to describe the RAID-6 data layout and parity computation.

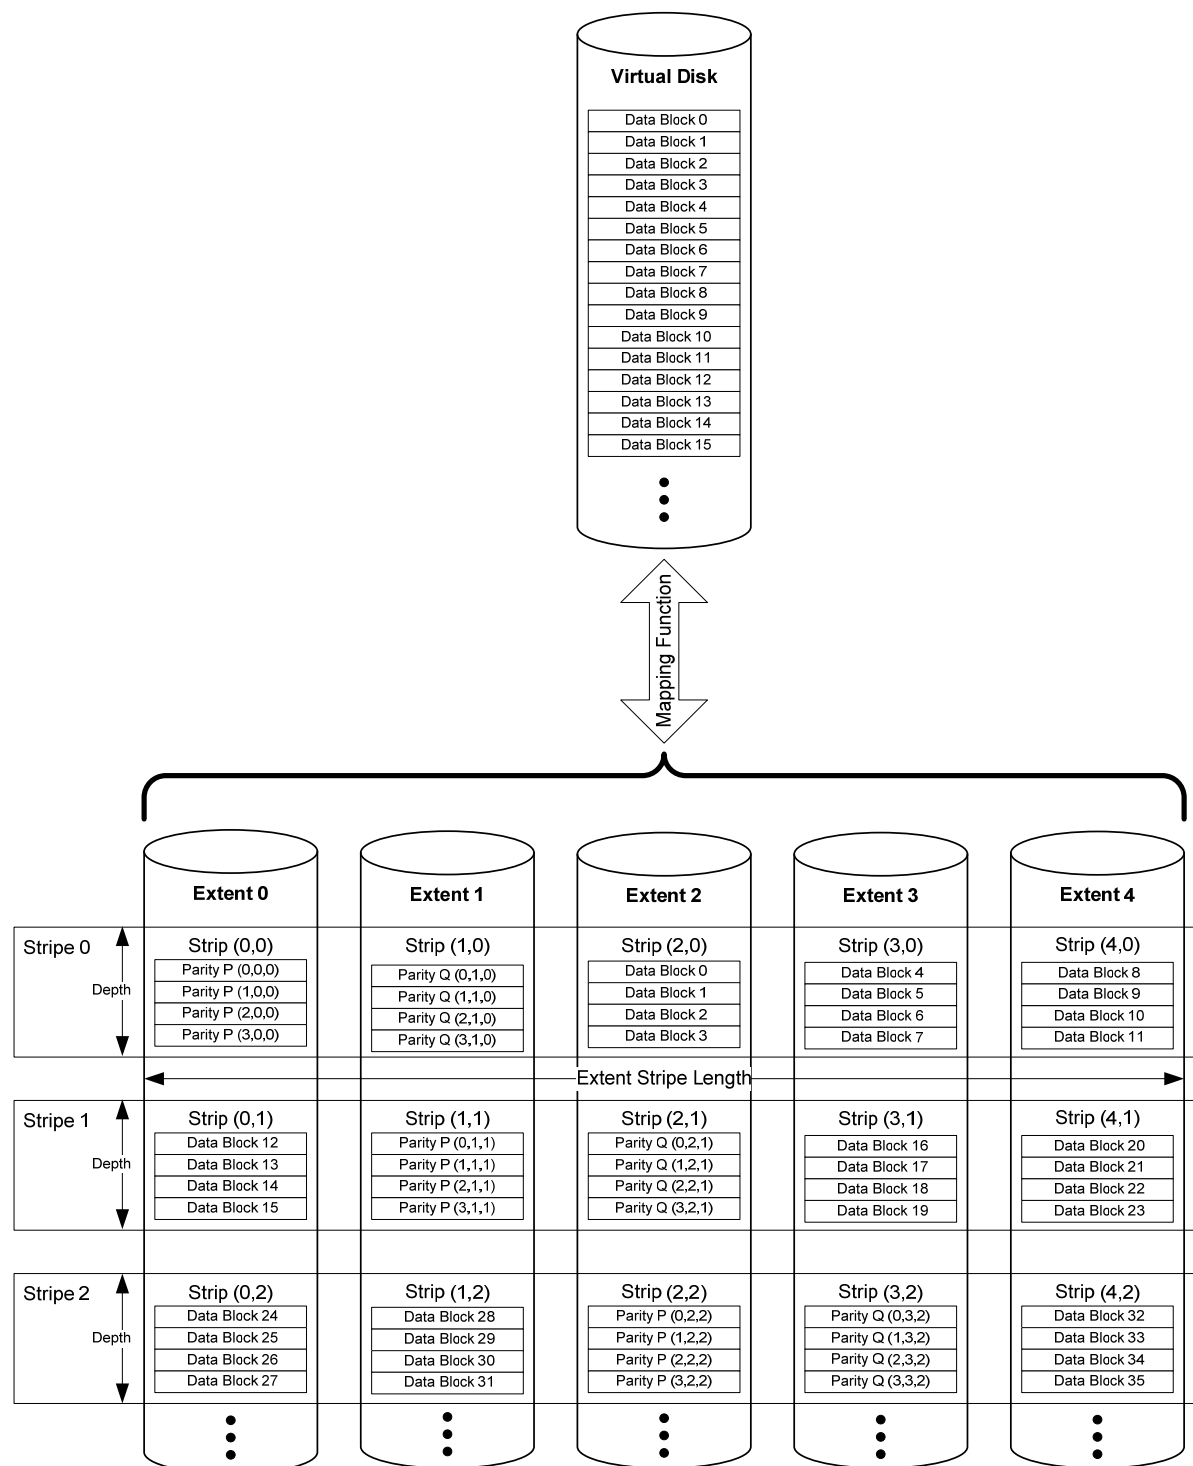


Figure 23: RAID-6 Rotating Parity 0 with Data Restart (PRL=06, RLQ=01) Example

Table 11: Indices and Constants for RAID-6 Rotating Parity 0 and N

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD	0	$N-1$
j	Index of a stripe in a VD	0	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/(N-2))$
k	Offset of a block from the beginning of a strip.	0	$L-1$
L	Size of a strip in blocks.	N/A (a fixed value)	N/A (a fixed value)
M	Number of data blocks in a VD. M MUST be evenly divisible by $(N-2)*L$.	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD.	N/A (a fixed value)	N/A (maximum of 255)
b	Index of a byte in an extent block.	0	Block_Size - 1
p	Index of the extent on which parity P for a given stripe resides.	0	$N-1$
q	Index of the extent on which parity Q for a given byte stripe resides.	0	$N-1$
Z	The polynomial used to generate the Galois field elements (required for parity computation).	N/A (a constant value)	N/A (a constant value)
K_i	The i^{th} Galois field element enumerated in Table 12.	See Table 12	See Table 12

The stripe j on which a virtual block x resides MUST adhere to the following formula:

Eq. 71

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-2))$$

The extent p on which the parity P for a given stripe j resides MUST adhere to the following formula:

Eq. 72

$$p = \text{MOD} (\text{FLOOR} ((\text{FLOOR} (x/L))/(N-2)), N)$$

The extent q on which the parity Q for a given stripe j resides must adhere to the following formula:

Eq. 73

$$q = \text{MOD} (1+p, N)$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 74

```

IF MOD (FLOOR (x/L), N-2) < p
THEN {
    i = MOD (FLOOR (x/L), N-2)
    IF p+2 > N
    THEN i = i + MOD (p+2, N)
}
ELSE
    i = MOD (FLOOR (x/L), N-2) + 2

```

The allocation of data blocks in a Rotating Parity 0 with Data Restart VD MUST adhere to the following formula:

Eq. 75

$$\text{virtual_block} (x) = \text{extent_block} (\text{MOD}(x, L), i, j)$$

The values of the parity P and Q MUST be computed according to the following formula:

Eq. 76

$$\text{Parity } P (k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p, i \neq q} \text{extent_block} (k, i, j).$$

The operator \oplus refers to bit-wise XOR of the operands.

Eq. 77

$$\text{Parity_Byte } Q (b, k, q, j) = \sum_{i=0}^{i=N-1, i \neq p, i \neq q} K_i \otimes \text{extent_block_byte} (b, k, i, j)$$

The operators \sum and \otimes refers to Galois field addition and Galois field multiplication respectively. The value for constant K_i MUST adhere to the following formula:

Eq. 78

$$K_i = \text{GFILOG}(i)$$

Eq. 79

$$\text{Parity Q}(k, q, j) = \bigvee_{b=0}^{b < \text{Block_Size}} \text{Parity_Byte Q}(b, k, q, j)$$

Notes:

1. The notation \bigvee implies that Parity Q (k, q, j) is a sequential enumeration of the Parity_Byte Q (b, k, q, j), where b ranges from 0 to the Block_Size -1.
2. Function extent_block_byte (b, k, i, j) returns the b^{th} byte of an extent_block (k, i, j).

4.2.22.1 Parity Re-computation on Block Update

If a data block is updated at byte index b of an extent u , the equations for P and Q parity re-computation MUST be as follows:

Eq. 80

$$\text{Parity P}_{\text{new}} = \text{extent_block_byte}(b, k, p, j) \oplus \text{old_extent_block_byte}(b, k, u, j) \oplus \text{new_extent_block_byte}(b, k, u, j)$$

Eq. 81

$$\text{Parity Q}_{\text{new}} = \text{extent_block_byte}(b, k, q, j) \oplus K_u \otimes \text{old_extent_block_byte}(b, k, u, j) \oplus K_u \otimes \text{new_extent_block_byte}(b, k, u, j)$$

Where $K_u = \text{GFILOG}(u)$ and \otimes refers to Galois multiplication operation. P_{new} and Q_{new} MUST be re-computed for every byte index b updated in the extent u .

4.2.22.2 Galois Field Operations

This section describes how the Galois field operations MUST be defined to calculate Parity Q for the RAID 6 algorithms described in Sections 4.2.22, 4.2.23, and 4.2.24.

4.2.22.2.1 GFILOG () Function (Z = 0x11D)

The following table lists the Galois field elements generated from the polynomial 0x11D.

Table 12: Galois Field Elements for Polynomial 0x11D

GFILOG(0xRS)		S															
		0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
R	0	1	2	4	8	10	20	40	80	1D	3A	74	E8	CD	87	13	26
	1	4C	98	2D	5A	B4	75	EA	C9	8F	03	06	0C	18	30	60	C0
	2	9D	27	4E	9C	25	4A	94	35	6A	D4	B5	77	EE	C1	9f	23
	3	46	8C	05	0A	14	28	50	A0	5D	BA	69	D2	B9	6F	DE	A1
	4	5F	BE	61	C2	99	2F	5E	BC	65	CA	89	0F	1E	3C	78	F0
	5	FD	E7	D3	BB	6B	D6	B1	7F	FE	E1	DF	A3	5B	B6	71	E2
	6	D9	AF	43	86	11	22	44	88	0D	1A	34	68	D0	BD	67	CE
	7	81	1F	3E	7C	F8	ED	C7	93	3B	76	EC	c5	97	33	66	CC
	8	85	17	2E	5C	B8	6D	DA	A9	4F	9E	21	42	84	15	2A	54
	9	A8	4D	9A	29	52	A4	55	AA	49	92	39	72	E4	D5	B7	73
	A	E6	D1	BF	63	C6	91	3F	7E	FC	E5	D7	B3	7B	F6	F1	FF
	B	E3	DB	AB	4B	96	31	62	C4	95	37	6E	DC	A5	57	AE	41
	C	82	19	32	64	C8	8D	07	0E	1C	38	70	E0	DD	A7	53	A6
	D	51	A2	59	B2	79	F2	F9	EF	C3	9B	2B	56	AC	45	8A	09
	E	12	24	48	90	3D	7A	F4	F5	F7	F3	FB	EB	CB	8B	0B	16
	F	2C	58	B0	7D	FA	E9	CF	83	1B	36	6C	D8	AD	47	8E	XX

Note: The values within Table 12 are hexadecimal values.

GFILOG(0xFF) is undefined and represented as XX. Here are two examples of how to apply Table 12:

$$\text{GFILOG}(0x8C) = 0x84,$$

$$\text{GFILOG}(0x21) = 0x27.$$

4.2.22.2.2 GFLOG() Function (Z = 0x11D)

The following table lists the logarithmic value of the Galois field elements.

Table 13: Log of Galois Field Elements for Polynomial 0x11D

GFLOG(0xRS)		S															
		0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
R	0	XX	00	01	19	02	32	1A	C6	03	DF	33	EE	1B	68	C7	4B
	1	04	64	E0	0E	34	8D	EF	81	1C	C1	69	F8	C8	08	4C	71
	2	05	8A	65	2F	E1	24	0F	21	35	93	8E	DA	F0	12	82	45
	3	1D	B5	C2	7D	6A	27	F9	B9	C9	9A	09	78	4D	E4	72	A6
	4	06	BF	8B	62	66	DD	30	FD	E2	98	25	B3	10	91	22	88
	5	36	D0	94	CE	8F	96	DB	BD	F1	D2	13	5C	83	38	46	40
	6	1E	42	B6	A3	C3	48	7E	6E	6B	3A	28	54	FA	85	BA	3D
	7	CA	5E	9B	9F	0A	15	79	2B	4E	D4	E5	AC	73	F3	A7	57
	8	07	70	C0	F7	8C	80	63	0D	67	4A	DE	ED	31	C5	FE	18
	9	E3	A5	99	77	26	B8	B4	7C	11	44	92	D9	23	20	89	2E
	A	37	3F	D1	5B	95	BC	CF	CD	90	87	97	B2	DC	FC	BE	61
	B	F2	56	D3	AB	14	2A	5D	9E	84	3C	39	53	47	6D	41	A2
	C	1F	2D	43	D8	B7	7B	A4	76	C4	17	49	EC	7F	0C	6F	F6
	D	6C	A1	3B	52	29	9D	55	AA	FB	60	86	B1	BB	CC	3E	5A
	E	CB	59	5F	B0	9C	A9	A0	51	0B	F5	16	EB	7A	75	2C	D7
	F	4F	AE	D5	E9	E6	E7	AD	E8	74	D6	F4	EA	A8	50	58	AF

Note: The values within Table 13 are hexadecimal values.

GFLOG(0x00) is undefined and represented as XX. Here are three examples of how to apply Table 13.

$$\text{GFLOG}(0x01) = 0x00,$$

$$\text{GFLOG}(0xF9) = 0xD6,$$

$$\text{GFLOG}(0x94) = 0x26.$$

4.2.22.2.3 Galois Field Addition (\oplus)

The Galois field addition operation is defined as the bitwise XOR (\oplus) of the operands. For example,

Eq. 82

$$\Sigma (0x22, 0x33, 0x44) = (0x22 \oplus 0x33 \oplus 0x44) = 0x55.$$

4.2.22.2.4 Galois Field Multiplication (\otimes)

The Galois field multiplication ($a \otimes b$) is defined as:

Eq. 83

$$\text{GFLOG}(\text{MOD}(\text{GFLOG}(a) + \text{GFLOG}(b), 0xFF)).$$

Note: Integer addition applies within Eq. 83.

Refer Section 4.2.22.2 and Section 4.2.22.1 for the definitions of functions GFLOG() and GFILOG() respectively. For example:

1. $0x33 \otimes 0x44 = \text{GFILOG}(\text{MOD}(0x7D + 0x66, 0xFF)) = \text{GFILOG}(0xE3) = 0x90$
2. $0x08 \otimes 0x54 = \text{GFILOG}(\text{MOD}(0x03 + 0x8F, 0xFF)) = \text{GFILOG}(0x92) = 0x9A$

4.2.23 RAID-6 Rotating Parity N with Data Restart (PRL=06, RLQ=02)

Figure 24 gives an example of an implementation of RAID-6 called Rotating Parity N with Data Restart. The indices and constants provided in Table 11 are valid for this type of RAID.

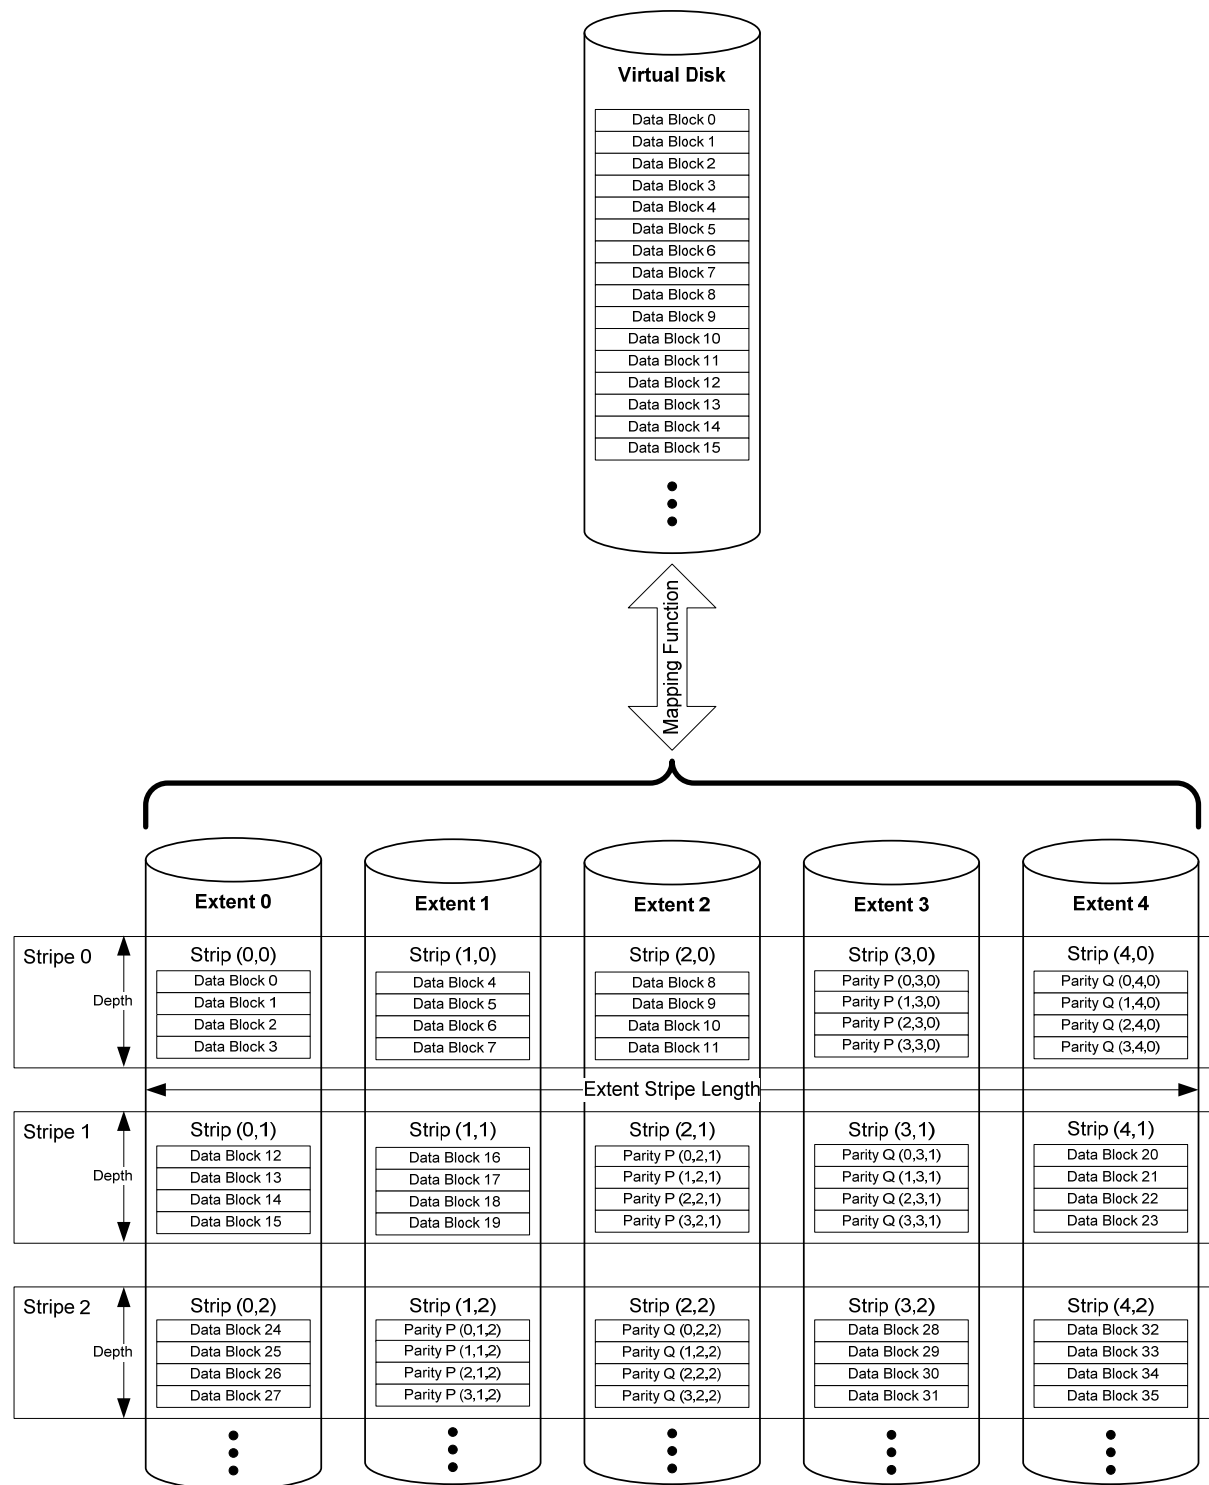


Figure 24: RAID-6 Rotating Parity N with Data Restart (PRL=06, RLQ=02) Example

The stripe j on which a virtual block x resides MUST adhere to the following formula:

Eq. 84

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-2))$$

The extent p on which the parity P for a given stripe j resides must adhere to the following formula:

Eq. 85

$$p = (N-1) - (\text{MOD}(\text{FLOOR}(\text{FLOOR}(x/L)/(N-2) + 1), N))$$

The extent q on which the parity Q for a given stripe j resides must adhere to the following formula:

Eq. 86

$$q = \text{MOD}(1+p, N)$$

The extent i on which a given block x resides MUST adhere to the following formula:

Eq. 87

```

IF MOD (FLOOR (x/L), N-2) < p
THEN {
    i = MOD (FLOOR (x/L), N-2)
    IF p+2 > N
    THEN i = i + MOD (p+2, N)
}
ELSE
    i = MOD (FLOOR (x/L), N-2) + 2

```

The allocation of data blocks in a Rotating Parity N with Data Restart VD MUST adhere to the following formula:

Eq. 88

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), i, j)$$

The values of the parity P and Q MUST be computed according to the following formula:

Eq. 89

$$\text{Parity P}(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p, i \neq q} \text{extent_block}(k, i, j)$$

and the operator \oplus refers to bit-wise XOR of the operands.

Eq. 90

$$\text{Parity_Byte Q } (b,k,q,j) = \sum_{i=0}^{i=N-1, i \neq p, i \neq q} K_i \otimes \text{extent_block_byte } (b,k,i,j)$$

The operators \sum and \otimes refer to Galois field addition and Galois field multiplication respectively. The value for constant K_i is arrived as:

Eq. 91

$$K_i = \text{GFILOG}(i).$$

Eq. 92

$$\text{Parity Q } (k,q,j) = \bigvee_{b=0}^{b < \text{Block_Size}} \text{Parity_Byte Q } (b,k,q,j)$$

Notes:

1. The notation \bigvee implies that Parity Q (k,q,j) is a sequential enumeration of the Parity_Byte Q (b,k,q,j) , where b ranges from 0 to the Block_Size - 1.
2. Function extent_block_byte (b,k,i,j) returns the b^{th} byte of an extent_block (k,i,j) .

4.2.23.1 Parity Re-computation on Block Update

If a data block is updated at byte index b of an extent u , the equations for P and Q parity re-computation MUST be as follows:

Eq. 93

$$\text{Parity P}_{\text{new}} = \text{extent_block_byte}(b,k,p,j) \oplus \text{old_extent_block_byte}(b,k,u,j) \oplus \text{new_extent_block_byte}(b,k,u,j)$$

Eq. 94

$$\text{Parity Q}_{\text{new}} = \text{extent_block_byte}(b,k,q,j) \oplus K_u \otimes \text{old_extent_block_byte}(b,k,u,j) \oplus K_u \otimes \text{new_extent_block_byte}(b,k,u,j)$$

Where $K_u = \text{GFILOG}(u)$ and \otimes refers to Galois multiplication operation. P_{new} and Q_{new} MUST be re-computed for every byte index b updated in the extent u .

4.2.23.2 Galois Field Operations

The Galois field operations that MUST be used to calculate Parity Q for the RAID 6 algorithm described in Section 4.2.23 are the same as the Galois field operations described in 4.2.22.2.

4.2.24 RAID 6 Rotating Parity N with Data Continuation (PRL=06, RLQ=03)

Figure 25 gives an example of an implementation of RAID-6 called Rotating Parity N with Data Continuation. The indices and constants provided in Table 11 are valid for this type of RAID.

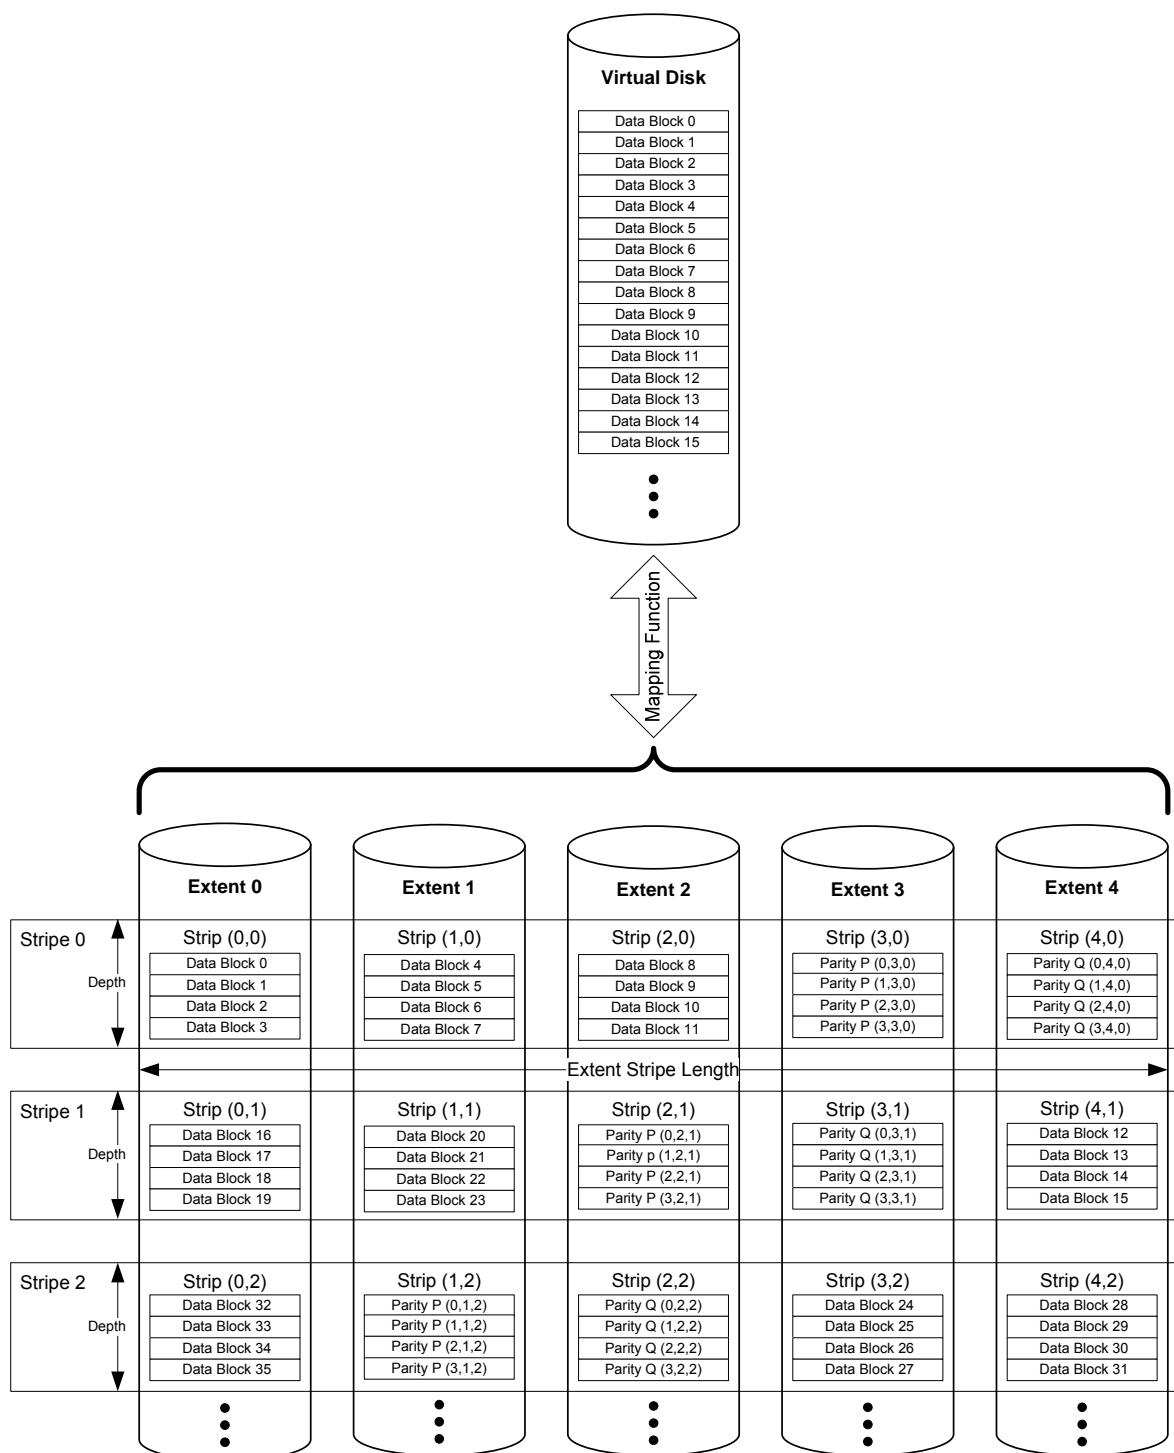


Figure 25: RAID-6 Rotating Parity N with Data Continuation (PRL=06, RLQ=03) Example

The stripe j on which a virtual block x resides MUST adhere to the following formula:

Eq. 95

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-2))$$

The extent p on which the parity P for a given stripe j resides must adhere to the following formula:

Eq. 96

$$p = (N-1) - (\text{MOD}(\text{FLOOR}(\text{FLOOR}(x/L)/(N-2) + 1), N))$$

The extent q on which the parity Q for a given stripe j resides must adhere to the following formula:

Eq. 97

$$q = \text{MOD}(1+p, N)$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 98

$$i = \text{MOD}(\text{MOD}(\text{FLOOR}(x/L), N-2) + p + 1, N).$$

The allocation of data blocks in a RAID-6 Rotating Parity N with Data Restart VD MUST adhere to the following formula:

Eq. 99

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), i, j).$$

The values of the parity P and Q MUST be computed according to the following formula:

Eq. 100

$$\text{Parity } P(k, p, j) = \bigoplus_{i=0}^{N-1, i \neq p, i \neq q} \text{extent_block}(k, i, j)$$

and the operator \oplus refers to bit-wise XOR of the operands.

Eq. 101

$$\text{Parity_Byte } Q(b, k, q, j) = \sum_{i=0}^{i=N-1, i \neq p, i \neq q} K_i \otimes \text{extent_block_byte}(b, k, i, j)$$

The operators \sum and \otimes refer to Galois field addition and Galois field multiplication respectively. The value for constant K_i is arrived as:

Eq. 102

$$K_i = \text{GFILOG}(i).$$

Eq. 103

$$\text{Parity } Q(k, q, j) = \bigvee_{b=0}^{b < \text{Block_Size}} \text{Parity_Byte } Q(b, k, q, j)$$

Notes:

1. The notation \bigvee implies that Parity $Q(k, q, j)$ is a sequential enumeration of the Parity_Byte $Q(b, k, q, j)$, where b ranges from 0 to the Block_Size - 1.
2. Function `extent_block_byte(b, k, i, j)` returns the b^{th} byte of an extent_block (k, i, j) .

4.2.24.1 Parity Re-computation on Block Update

If a data block is updated at byte index b of an extent u , the equations for P and Q parity re-computation MUST be as follows:

Eq. 104

$$\text{Parity } P_{\text{new}} = \text{extent_block_byte}(b, k, p, j) \oplus \text{old_extent_block_byte}(b, k, u, j) \oplus \text{new_extent_block_byte}(b, k, u, j)$$

Eq. 105

$$\text{Parity } Q_{\text{new}} = \text{extent_block_byte}(b, k, q, j) \oplus K_u \otimes \text{old_extent_block_byte}(b, k, u, j) \oplus K_u \otimes \text{new_extent_block_byte}(b, k, u, j)$$

Where $K_u = \text{GFLOG}(u)$ and \otimes refers to Galois multiplication operation. P_{new} and Q_{new} MUST be re-computed for every byte index b updated in the extent u .

4.2.24.2 Galois Field Operations

The Galois field operations that MUST be used to calculate Parity Q for the RAID 6 algorithm described in Section 4.2.24 are the same as the Galois field operations described in 4.2.22.2.

4.2.25 Multi Disk Failure RAID Rotating Parity 0 with Data Restart (PRL=07, RLQ=00)

Figure 26 gives an example of Multi Disk Failure (MDF) RAID Rotating Parity 0 with Data Restart. This example supports two simultaneous disk failures.

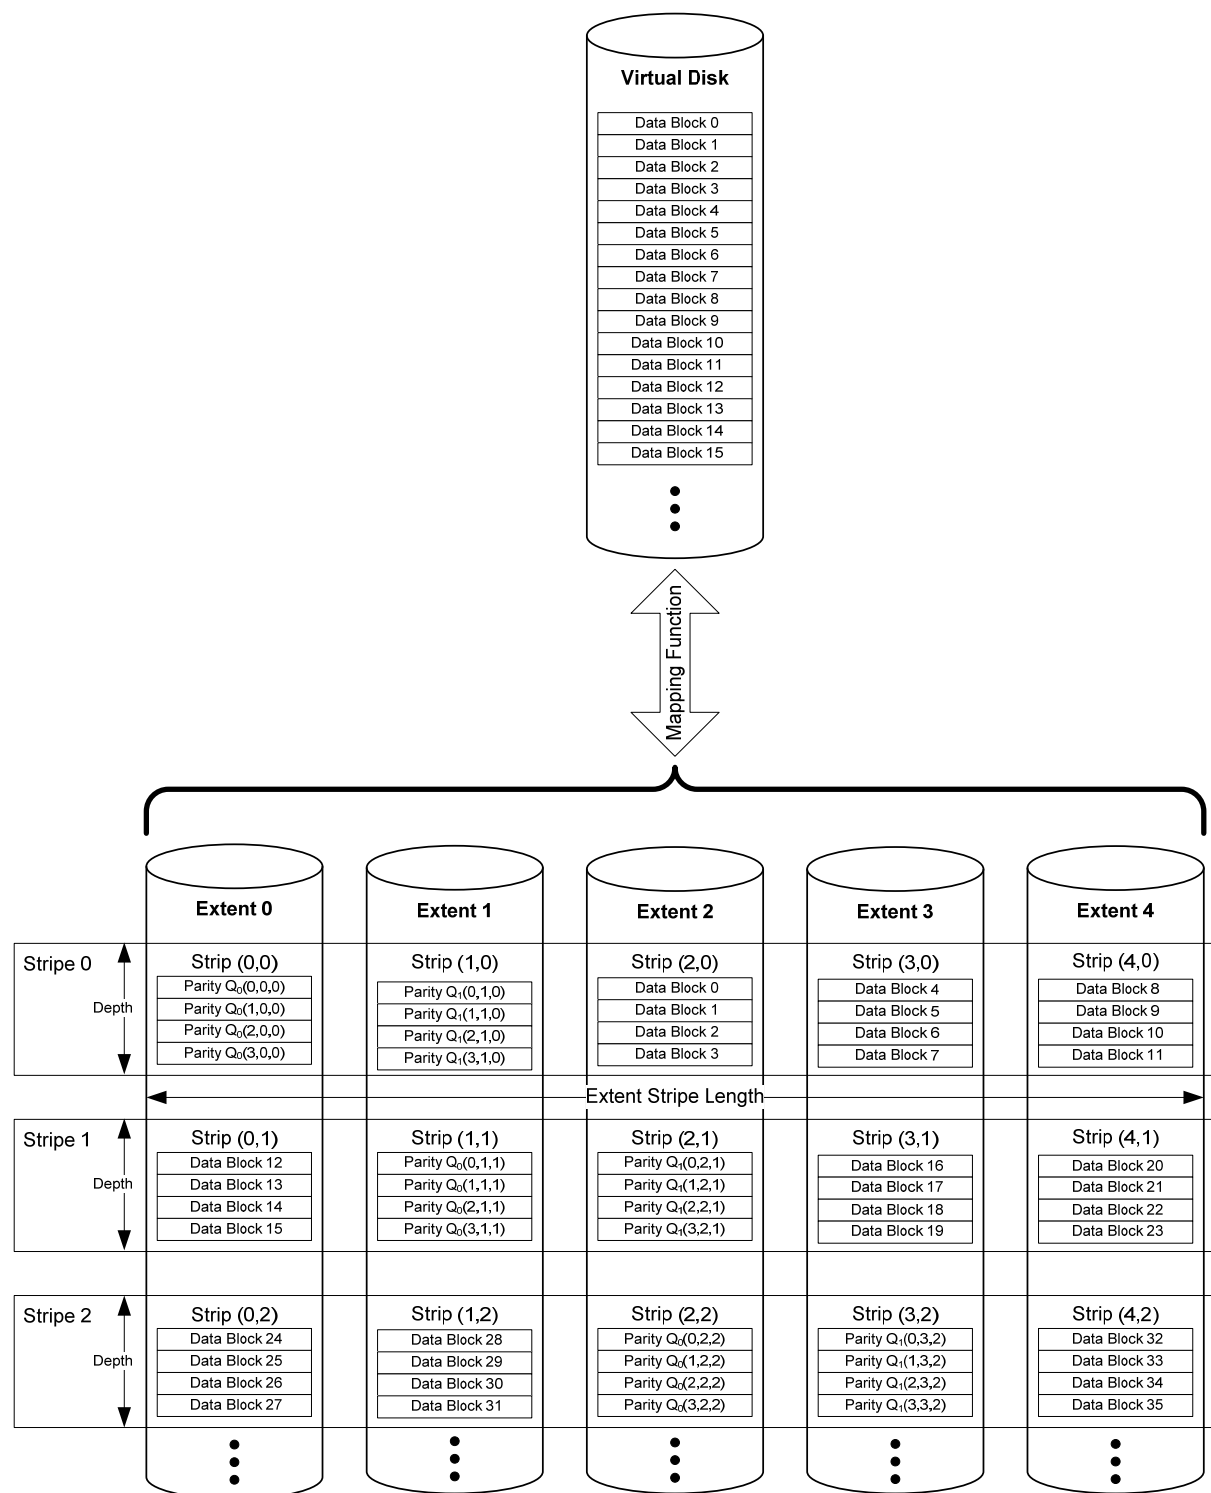


Figure 26: MDF RAID Rotating Parity 0 with Data Restart (PRL=07, RLQ=00) Example

Table 14 defines the constants and indices that are required to describe the Multi Disk Failure (MDF) RAID data layout and parity computation.

Table 14: Indices and Constants for Multi Disk Failure RAID

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of an extent of a VD.	0	$N-1$
j	Index of a stripe in a VD.	0	$\text{FLOOR}(\text{FLOOR}((M-1)/L)/(N-F))$
k	Offset of a block from the beginning of a strip.	0	$L-1$
L	Size of a strip in blocks.	N/A (a fixed value)	N/A (a fixed value)
M	Number of data blocks in a VD. M MUST be evenly divisible by $(N-F)*L*R$.	N/A (a fixed value)	N/A (a fixed value)
N	Number of extents in a VD.	N/A (a fixed value)	N/A (maximum of 255)
b	Index of a byte in an extent block.	0	Block_Size - 1.
F	Maximum number of simultaneous disk failures tolerated.	2	127
f	Index to the Parity Q_f . Relative index within the Parity extents.	0	$F-1$
Q_f	The f^{th} Parity Block.	N/A	N/A
q_f	Index of the extent on which f^{th} parity Q_f for a given stripe resides.	0	$N+F-1$
R	Number of stripes before rotating Parity (where $R=2^n$ and n is a positive integer).	1	0x80000000

The stripe j on which a virtual block x resides MUST adhere to the following formula:

Eq. 106

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-F))$$

The extent q_f on which the f^{th} parity Q_f for a given stripe j resides MUST adhere to the following formulas:

Eq. 107

$$q_0 = \text{MOD} (\text{FLOOR}(\text{FLOOR}(x/(L * R)) / (N-F)), N), \text{ when } f=0;$$

Eq. 108

```
For (f=1; f ≤ F-1; f++) {  
    qf = MOD (q(f-1)+1, N)  
}
```

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 109

```
IF MOD (FLOOR (x/(L * R)), N-F) < q0  
THEN {  
    i = MOD (FLOOR (x/(L * R)), N-F)  
    IF q0+F > N  
    THEN i = i + MOD (q0+F, N)  
}  
ELSE  
    i = MOD (FLOOR (x/(L * R)), N-F) + F
```

The allocation of data blocks in a MDF RAID Rotating Parity 0 with Data Restart VD MUST adhere to the following formula:

Eq. 110

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), i, j)$$

4.2.25.1 Galois Field Generation

The Galois field elements are used as the basis in the computation of parity bytes Q_i , and are dependent on the Generator Polynomial. The Galois field elements MUST be generated using the following algorithm.

Eq. 111

```
int GenerateGaloisField (int nFieldSize, unsigned int IPolynomial)  
{  
    unsigned int    b, index, gf_elements=1;  
    unsigned short * GFLOG, * GFILOG;  
  
    gf_elements <= 256; // nFieldSize;  
  
    GFLOG = (unsigned short *) malloc (sizeof (unsigned short) * gf_elements);
```

```

GFILOG = (unsigned short *) malloc (sizeof (unsigned short) * gf_elements);

// Initialize the Galois Field and Galois-Log-Field Arrays respectively.
for (index=0; index <= gf_elements-1; index++)
    GFLOG[index]=GFILOG[index]=0x00;
b = 1;
for (index = 0; index <= gf_elements-1; index++) {
    GFILOG[index] = (unsigned char) b;          // Galois Field Element.
    GFLOG[b] = (unsigned char) index;          // Log of Galois Field Element.
    b ^= ((b<=1) & gf_elements) ? IPolynomial : 0;
}
return 0;
}

```

The value of the generator polynomial, *IPolynomial*, MUST be one of the polynomials defined in the MDF Parity Generator Polynomial field of the Virtual Disk Configuration record (see 5.9.1).

4.2.25.2 Constant Matrix Generation

This section describes the generation of the Constant Matrix according to the Vandermonde Matrix Method^{1,2}. The elements of the Constant Matrix shall be used in the generation of Q_i parities and its use is described in 4.2.25.3.

Step 0:

Create the Vandermonde matrix (Figure 27) with the following properties:

- Any element in the Matrix (i, j) MUST be equal to i^j
- Number of rows MUST be equal to $N+F$
- Number of columns MUST be equal to N

$$\begin{pmatrix}
 0^0 & 0^1 & 0^2 & \dots & \dots & \dots & 0^{(N-1)} \\
 1^0 & 1^1 & 1^2 & \dots & \dots & \dots & 1^{(N-1)} \\
 2^0 & 2^1 & 2^2 & \dots & \dots & \dots & 2^{(N-1)} \\
 \dots & \dots & \dots & \dots & \dots & \dots & \dots \\
 (N-1)^0 & (N-1)^1 & (N-1)^2 & \dots & \dots & \dots & (N-1)^{(N-1)} \\
 N^0 & N^1 & N^2 & \dots & \dots & \dots & N^{(N-1)} \\
 \dots & \dots & \dots & \dots & \dots & \dots & \dots \\
 (N+F-1)^0 & (N+F-1)^1 & (N+F-1)^2 & \dots & \dots & \dots & (N+F-1)^{(N-1)}
 \end{pmatrix}$$

Figure 27 Vandermonde Matrix

¹ James S. Plank, "A Tutorial on Reed-Solomon Coding for Fault-Tolerance in RAID-like Systems," Software – Practice & Experience, Volume 27, Issue 9, September, 1997 pp. 995-1012.

² James S. Plank, Ying Ding, "Note: Correction to the 1997 Tutorial on Reed-Solomon Coding," Software: Practice & Experience, Volume 35, Issue 2, February 2005, pp. 189-194.

Step 1:

Compute the Galois Field Values, $v_{(i,j)}$, for each element (i,j) in the above Vandermonde matrix using Eq. 112.

Eq. 112

$$v_{(i,j)} = \left(\bigotimes_0^{(j-1)} i \right)$$

$$v_{(i,j)} = \left(GFLOG(MOD(+_0^{(j-1)} GFLOG(i), 0xFF)) \right)$$

Where GFLOG() and GFLOG() are defined in Eq. 111. The resulting Galois Field Element Matrix is shown in Figure 28.

$$\begin{pmatrix} 1 & 0 & 0 & \dots & \dots & \dots & 0 \\ v_{(1,0)} & v_{(1,1)} & v_{(1,2)} & \dots & \dots & \dots & v_{(1,N-1)} \\ v_{(2,0)} & v_{(2,1)} & v_{(2,2)} & \dots & \dots & \dots & v_{(2,N-1)} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ v_{(N-1,0)} & v_{(N-1,1)} & v_{(N-1,2)} & \dots & \dots & \dots & v_{(N-1,N-1)} \\ v_{(N,0)} & v_{(N,1)} & v_{(N,2)} & \dots & \dots & \dots & v_{(N,N-1)} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ v_{(N-F+1,0)} & v_{(N-F+1,1)} & v_{(N-F+1,2)} & \dots & \dots & \dots & v_{(N-F+1,N-1)} \end{pmatrix}$$

Figure 28: Galois Field Element Matrix

Step 2:

Transform the rows 0 to (N-1) of the Galois Field Element Matrix (Figure 28) to an Identity Matrix through a series of the following column operations:

- Any column C_i may be swapped with column C_j
- Any column C_i may be replaced with $C_i * c$, where $c \neq 0$
- Any column C_i may be replaced with $C_i = C_i + c * C_j$, where $i \neq j$ and $c \neq 0$.

The resulting Identity Matrix is referred to as the Parity Generator Constant Matrix (Figure 29).

$$\begin{pmatrix} 1 & 0 & 0 & \dots & \dots & \dots & 0 \\ 0 & 1 & 0 & \dots & \dots & \dots & 0 \\ 0 & 0 & 1 & \dots & \dots & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & \dots & \dots & 1 \\ K_{N,0} & K_{N,1} & K_{N,2} & \dots & \dots & \dots & K_{N,(N-1)} \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ K_{N+F-1,0} & K_{N+F-1,1} & K_{N+F-1,2} & \dots & \dots & \dots & K_{N+F-1,(N-1)} \end{pmatrix}$$

Figure 29: Parity Generator Constant Matrix

4.2.25.3 Galois Parity Computation

The value of Parity Q_f , where $0 \leq f \leq (F-1)$, is computed according to the following formula:

Eq. 113

$$\text{Parity_Byte } Q_f(b, k, f, j) = \sum_{i=0, i \neq q_0, \dots, i \neq q_{(F-1)}}^{i=N+F-1, i \neq q_0, \dots, i \neq q_{(F-1)}} K_{(N+f,i)} \otimes \text{extent_block_byte}(b, k, i, j)$$

The operators \sum and \otimes refers to Galois field addition and Galois field multiplication respectively.

Eq. 114

$$\text{Parity } Q_f(k, f, j) = \bigvee_{b=0}^{b < \text{Block_Size}} \text{Parity_Byte } Q_f(b, k, q, j)$$

Notes:

1. The notation \bigvee implies that Parity $Q_f(k, f, j)$ is a sequential enumeration of the Parity_Byte $Q_f(b, k, f, j)$, where b ranges from 0 to the Block_Size-1.
2. Function extent_block_byte(b, k, i, j) returns the b^{th} byte of an extent_block(k, i, j).

4.2.26 Multi Disk Failure RAID Rotating Parity N with Data Restart (PRL=07, RLQ=02)

Figure 30 gives an example of Multi Disk Failure (MDF) RAID Rotating Parity N with Data Restart. This example supports two simultaneous disk failures. The indices and constants provided in Table 14 are valid for this type of RAID.

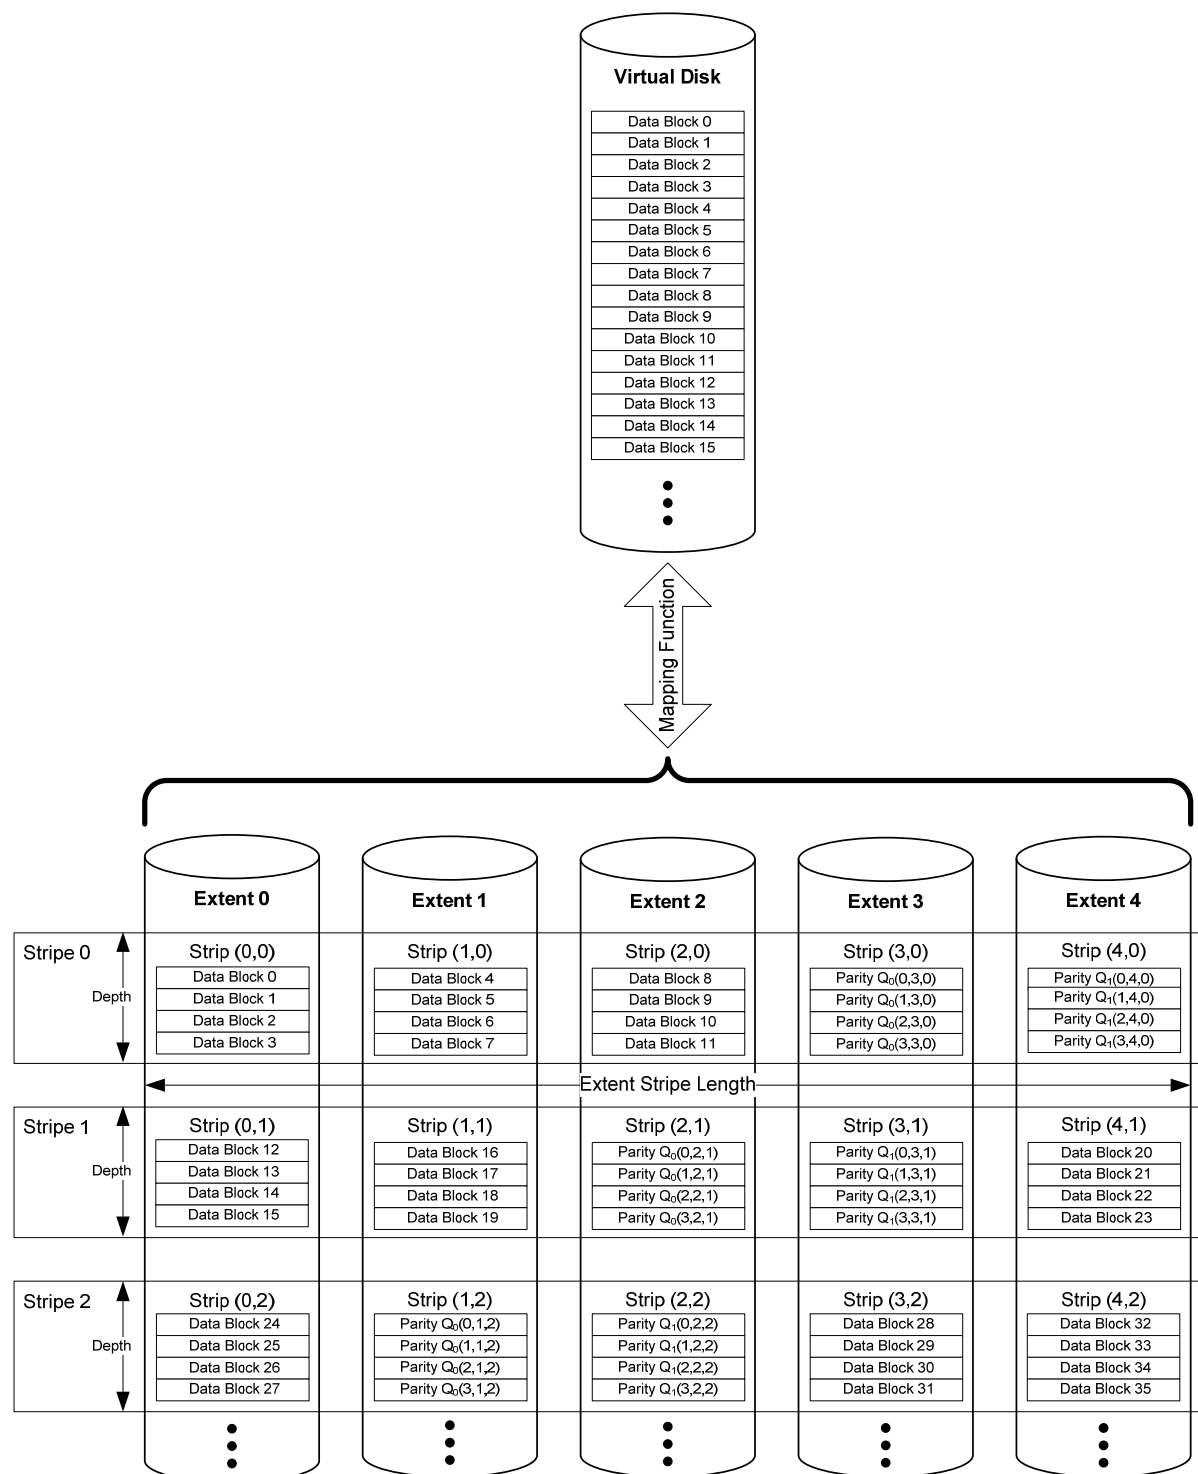


Figure 30: MDF RAID Rotating Parity N with Data Restart (PRL=07, RLQ=02) Example

The stripe j on which a virtual block x resides MUST adhere to the following formula:

Eq. 115

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-F))$$

The extent p_f on which the f^{th} parity Q_f for a given stripe j resides MUST adhere to the following formulas:

Eq. 116

$$q_0 = (N-1) - (\text{MOD}(\text{FLOOR}(\text{FLOOR}(x/(L \cdot R)) / (N-F) + (F-1)), N)), \text{ when } f=0;$$

Eq. 117

```
For (f=1; f ≤ F-1; f++) {
    q_f = MOD (q_{(f-1)}+f, N)
}
```

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 118

```
IF MOD (FLOOR (x/(L * R)), N-F) < q_0
THEN {
    i = MOD (FLOOR (x/(L * R)), N-F)
    IF q_0+F > N
    THEN i = i + MOD (q_0+F, N)
}
ELSE
    i = MOD (FLOOR (x/(L * R)), N-F) + F
```

The allocation of data blocks in a MDF RAID Rotating Parity N with Data Restart VD MUST adhere to the following formula:

Eq. 119

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), i, j)$$

The value of Parity Q_f is calculated as described in Sections 4.2.25.1, 4.2.25.2, and 4.2.25.3.

4.2.27 Multi Disk Failure RAID Rotating Parity N with Data Continuation (PRL=07, RLQ=03)

Figure 31 gives an example of Multi Disk Failure (MDF) RAID Rotating Parity N with Data continuation. This example supports two simultaneous disk failures. The indices and constants provided in Table 14 are valid for this type of RAID.

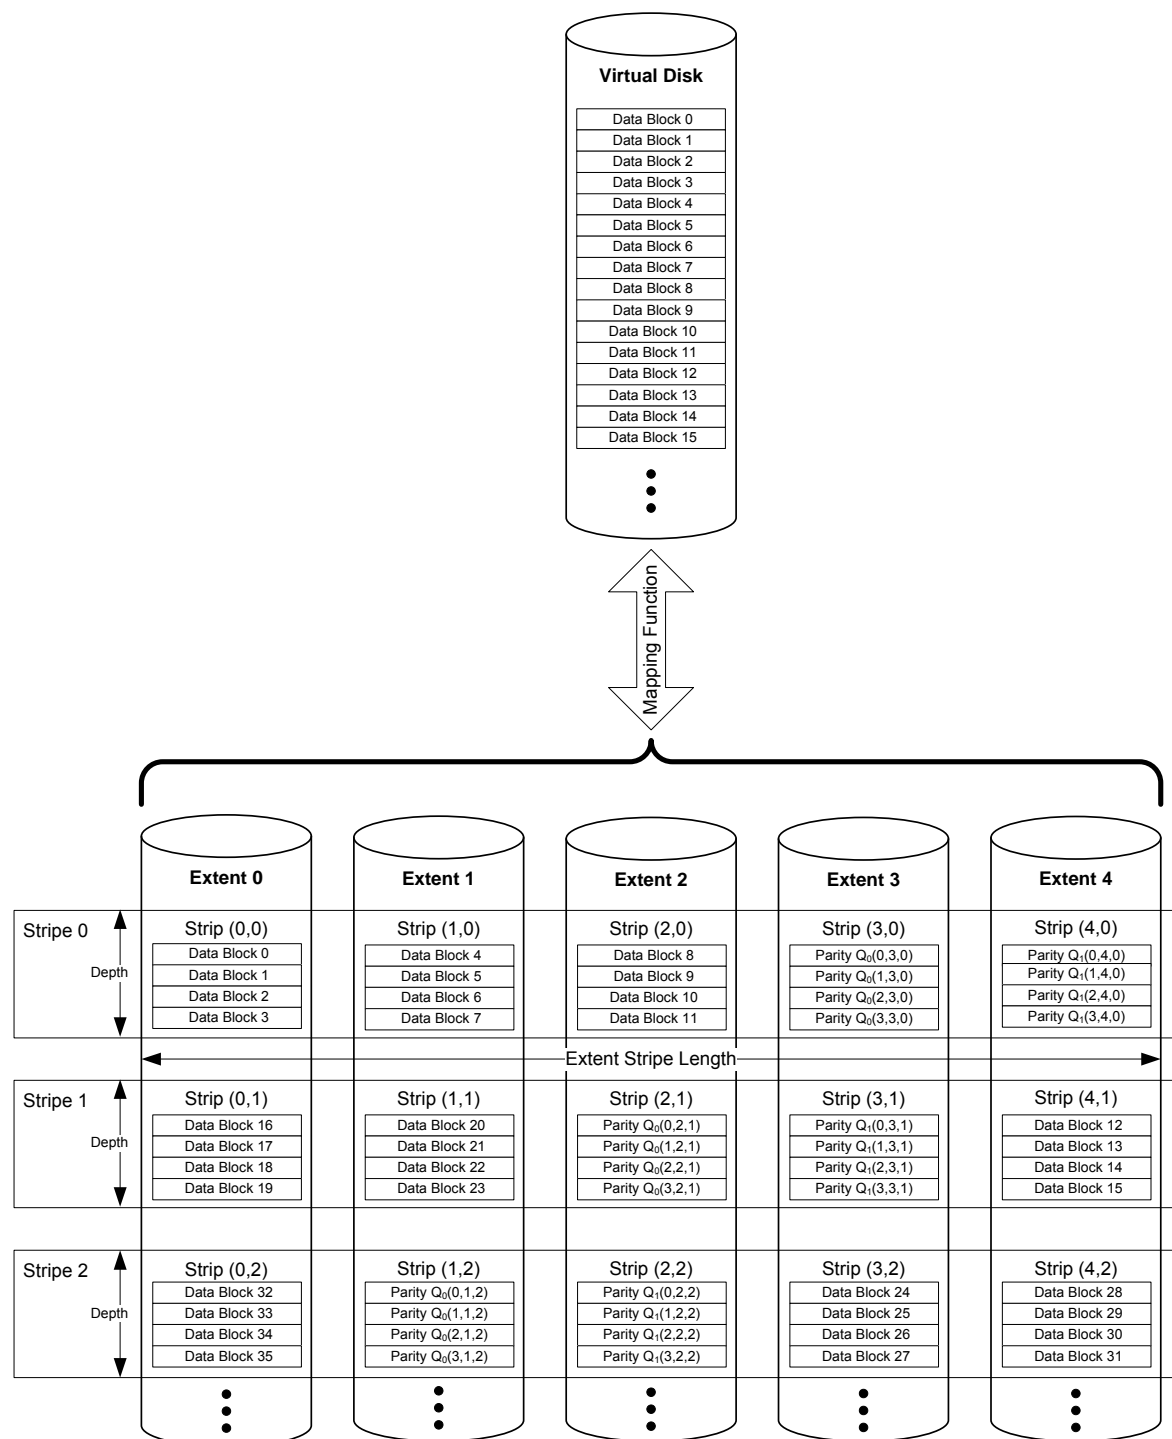


Figure 31: MDF RAID Rotating Parity N with Data Continuation (PRL=07, RLQ=03) Example

The stripe j on which a virtual block x resides MUST adhere to the following formula:

Eq. 120

$$j = \text{FLOOR}(\text{FLOOR}(x/L)/(N-F))$$

The extent q_f on which the f^{th} parity Q_f for a given stripe j resides MUST adhere to the following formulas:

Eq. 121

$$q_0 = (N-1) - (\text{MOD}(\text{FLOOR}(\text{FLOOR}(x/(L * R))) / (N-F) + (F-1)), N), \text{ when } f=0;$$

Eq. 122

$$\begin{aligned} &\text{For } (f=1; f \leq F-1; f++) \{ \\ &\quad q_f = \text{MOD}(q_{(f-1)} + f, N) \\ &\} \end{aligned}$$

The extent i on which a given virtual block x resides MUST adhere to the following formula:

Eq. 123

$$i = \text{MOD}(\text{MOD}(\text{FLOOR}(x/(L * R)), N-F) + q_0 + F, N).$$

The allocation of data blocks in a MDF RAID Rotating Parity N with Data Continuation VD MUST adhere to the following formula:

Eq. 124

$$\text{virtual_block}(x) = \text{extent_block}(\text{MOD}(x, L), i, j).$$

The value of Parity Q_f is calculated as described in Sections 4.2.25.1, 4.2.25.2, and 4.2.25.3.

4.3 Secondary RAID Level

Table 15 lists values used in the Secondary_RAID_Level field of the Virtual Disk Configuration Record (Section 5.9.1) and their definitions. The table defines secondary RAID levels such as Striped, Volume Concatenation, Spanned, and Mirrored for hybrid or multilevel virtual disks. The Secondary_RAID_Level field in the Virtual Disk Configuration Record MUST use the values defined in Table 15.

Table 15: Secondary RAID Levels

<u>Name</u>	<u>SRL Byte</u>	<u>Description</u>
Striped	0x00	Data is striped across Basic VDs. First strip stored on first BVD and next on next BVD. NOTE: BVD sequence is determined by the Secondary_Element_Seq field in the Virtual Disk Configuration Record (Section 5.9.1).
Mirrored	0x01	Data is mirrored across Basic VDs.
Concatenated	0x02	Basic VDs combined head to tail.
Spanned	0x03	A combination of stripping and concatenations involving Basic VDs of different sizes. NOTE: BVD sequence is determined by the Secondary_Element_Seq field in the Virtual Disk Configuration Record (Section 5.9.1).

4.3.1 Striped Secondary RAID Level (SRL=00)

In a VD with a Striped Secondary RAID Level, the data MUST be striped across basic virtual disks. The BVDs MUST have an equal number of user addressable data strips. The BVDs are not required to have the same Primary RAID Level (Section 4.1), RAID Level Qualifier (Section 4.2), or capacity. Table 16 gives the indices and constants used in the description of a Striped Secondary RAID Level (Striped SRL) VD.

Table 16: Indices and Constants for Striped Secondary RAID Level

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of a BVD of in a striped VD	0	$N-1$
j	Index of a stripe in a VD	0	$J-1$
q	Offset of a data block from the beginning of a stripe	0	$S-1$
L_i	Size of a strip in blocks on BVD i	Equal to the stripe depth of BVD i 's Primary RAID Level	Equal to the stripe depth of a BVD i 's Primary RAID Level
M	Number of blocks in a VD	N/A (a fixed value)	N/A (a fixed value)
N	Number of BVDs in a striped VD	N/A (a fixed value)	N/A (a fixed value)
J	Number of stripes	M/S	M/S
S	Size of a stripe in a striped VD	$S = \sum_{i=0}^{N-1} L_i$	$S = \sum_{i=0}^{N-1} L_i$
x	Offset of a data block from the beginning of a VD	0	$M-1$

The Striped Secondary RAID level introduces the concept of a *stripe_block*. Since BVDs in a Striped SRL VD can have different strip sizes, it is more economical to represent a block's offset from the beginning of a stripe rather than from the beginning of a strip on a specific BVD. A stripe_block uses the indices q and j to identify its position (i.e., stripe_block(q,j))

The size of a stripe on a Striped SRL VD MUST adhere to the following formula:

Eq. 125

$$S = \sum_{i=0}^{N-1} L_i$$

The number of stripes J is equal to M/S . The user addressable data MUST be distributed according to the following formula:

Eq. 126

$$\text{virtual_block}(x) = \text{stripe_block}(\text{MOD}(x, S), \text{FLOOR}(x, S))$$

Figure 32 gives an example of a VD with a Striped Secondary RAID Level. In this example $L_0 = 4$, $L_1 = 2$, $L_2 = 3$, $L_3 = 4$, and $L_4 = 3$. Thus, $S = 16$.

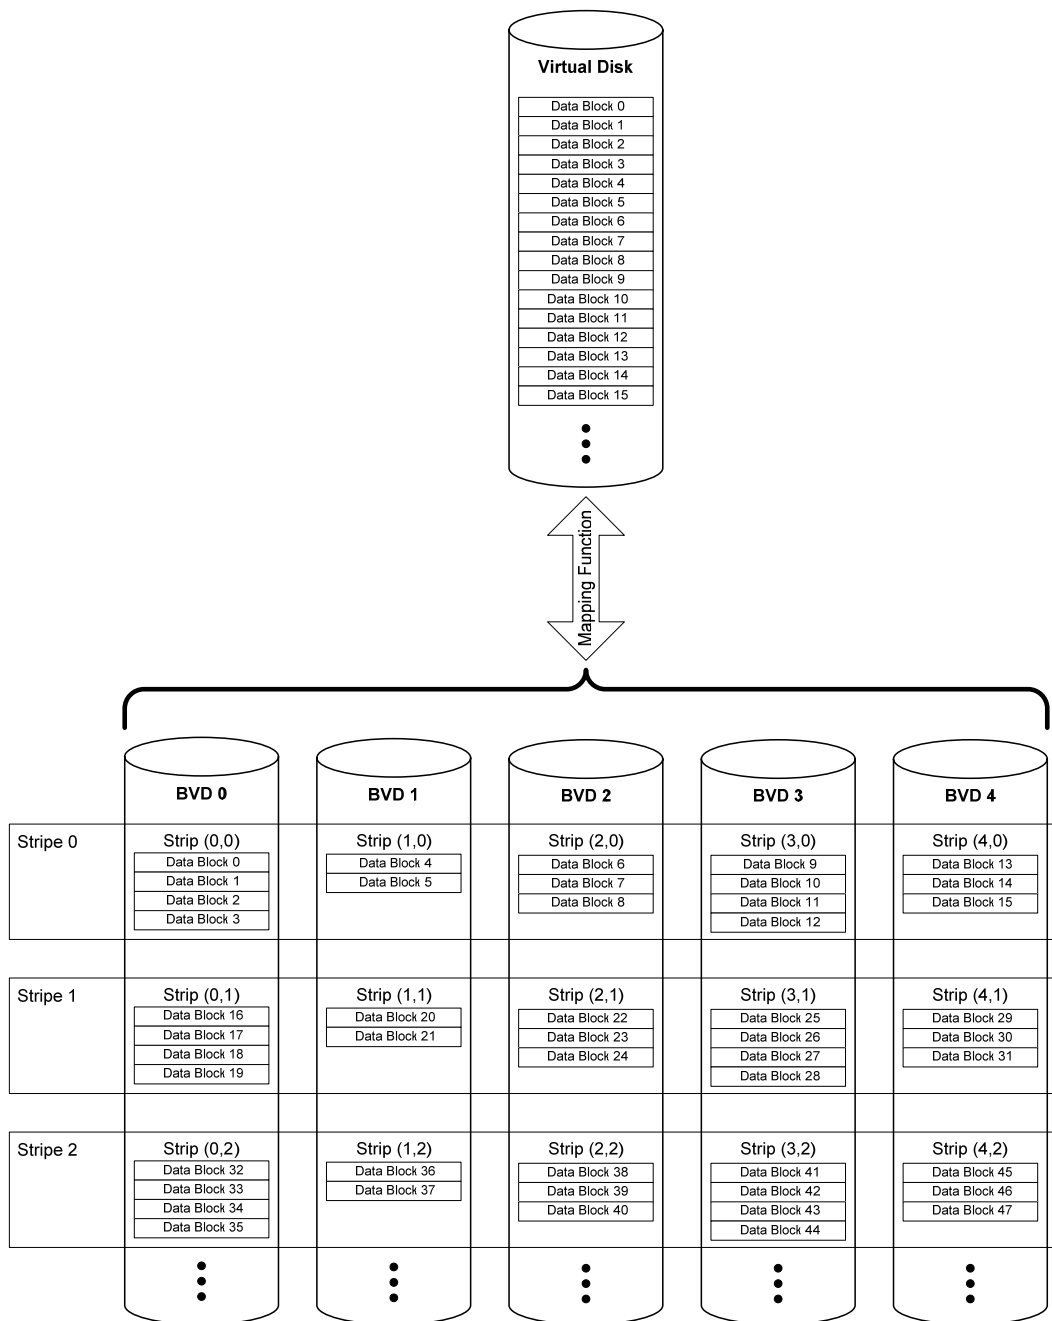


Figure 32: Striped Secondary RAID Level (SRL=00)

4.3.2 Mirrored Secondary RAID Level (SRL=01)

In a VD with a Mirrored Secondary RAID Level, the data **MUST** be mirrored across basic virtual disks. Figure 33 gives an example of a VD with a Mirrored Secondary RAID Level. Each data block of the virtual disk (virtual_block(x)) **MUST** be stored at the same user addressable data offset in each BVD.




Figure 33: Mirrored Secondary RAID Level (SRL=01)

4.3.3 Concatenated Secondary RAID Level (SLR=02)

In a VD with a Concatenated Secondary RAID Level, the user addressable data space is created by combining the user addressable data spaces of the BVDs making up the VD. The BVDs act as extents for the VD and are called *BVD extents*. The BVDs are combined in a head to tail fashion. The BVDs are not required to have the same addressable data space.

Figure 34 gives an example of a Concatenated VD with three BVD extents. Table 17 gives the indices and constants used in the formulas below

To represent a specific block in a specific BVD the following notation is used:

$\text{bvd_extent_block}(y, i)$.

The allocation of data MUST adhere to the following equations.

Eq. 127

```

IF  $x = 0$ 
THEN  $\text{virtual\_block}(0) = \text{bvd\_extent\_block}(0, 0)$ 
ELSE IF  $x = M-1$ 
THEN  $\text{virtual\_block}(x) = \text{bvd\_extent\_block}(C_{N-1}, N-1)$ 
ELSE IF  $\text{virtual\_block}(x-1) = \text{bvd\_extent\_block}(C_i-1, i)$ 
THEN  $\text{virtual\_block}(x) = \text{bvd\_extent\_block}(0, i+1)$ 
ELSE IF  $\text{virtual\_block}(x-1) = \text{bvd\_extent\_block}(y-1, i)$ 
THEN  $\text{virtual\_block}(x) = \text{bvd\_extent\_block}(y, i)$ 

```

Table 17: Indices and Constants for Concatenated Secondary RAID Level VDs

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of a BVD extent of a VD	0	$N-1$
M	Number of data blocks in a VD	N/A (a fixed value)	N/A (a fixed value)
N	Number of BVDs acting as extents in a VD with a Concatenated Secondary RAID Level	N/A (a fixed value)	N/A (a fixed value)
C_i	Number of data blocks in BVD i	N/A (a fixed value)	N/A (a fixed value)
x	Offset of a data block from the beginning of the VD	0	$M-1$
y	Offset of a data block from the beginning of a BVD extent	0	$M-1$

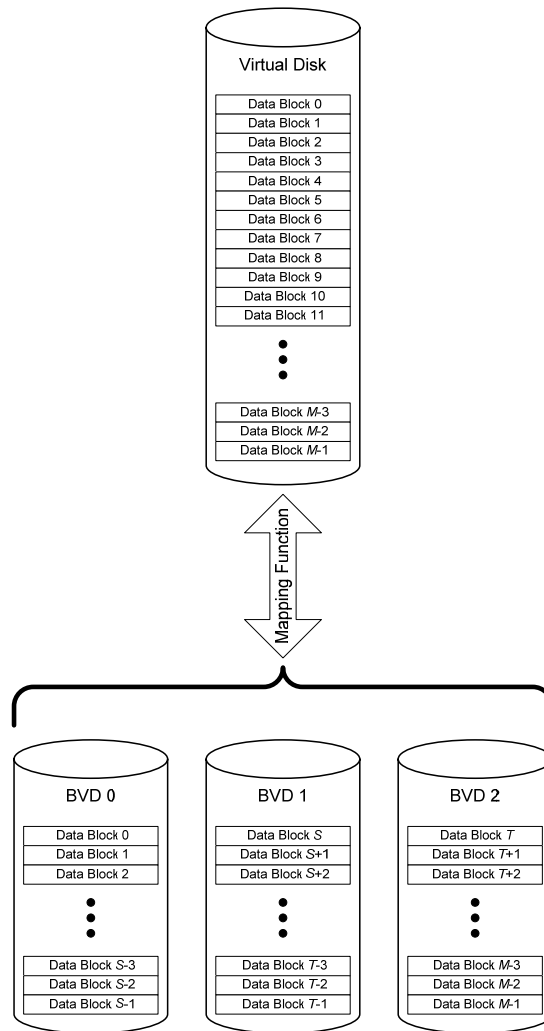


Figure 34: Concatenated Secondary RAID Level (SRL=02)

4.3.4 Spanned Secondary RAID Level (SRL=03)

In VD with a Spanned Secondary RAID Level, the user addressable data space is created by a combination of stripping and concatenation of the user addressable data spaces of the BVDs making up the VD. The BVDs can be in any order. They can have different capacities, different RAID levels and different strip sizes. Unlike the Striped Secondary RAID Level, the BVDs do not have to have the same number of user addressable data strips. Table 18 gives the indices and constants used in the description of a Spanned Secondary RAID Level (Spanned SRL) VD.

The size of stripe j MUST be determined using the following formula:

$$\text{Eq. } S_j = \sum_{i=0}^{N-1} \{ \text{IF } J_i \geq j \text{ THEN } L_i \text{ ELSE } 0 \}.$$

The total number of user addressable data blocks included in stripes 0 through j is:

Eq. 128

$$T_j = \sum_{v=0}^j S_v$$

The Spanned SRL VD is also described using the `stripe_block(q, j)` concept (Section 4.3.1). The user addressable data in a Spanned SRL VD MUST be distributed according to the following formula.

Eq. 129

$$\text{virtual_block}(x) = \text{stripe_block}(x - T_u, u + 1),$$

where T_u is the largest value of T_j less than or equal to x .

Table 18: Indices and Constants for Spanned Secondary RAID Level

Variable or Index	Description	Minimum Value	Maximum Value
i	Index of a BVD of in a spanned VD	0	$N-1$
j	Index of a stripe in a VD	0	$J-1$
q	Offset of a data block from the beginning of a stripe	0	$S-1$
L_i	Size of a strip in blocks on BVD i	Equal to the stripe depth of BVD i 's Primary RAID Level	Equal to the stripe depth of a BVD i 's Primary RAID Level
M	Number of blocks in a VD	N/A (a fixed value)	N/A (a fixed value)
N	Number of BVDs in a striped VD	N/A (a fixed value)	N/A (a fixed value)
J_i	Number of stripes in on BVD i	N/A (a fixed value determined by BVD i 's Primary RAID Level parameters)	N/A (a fixed value determined by BVD i 's Primary RAID Level parameters)
S_j	Size of a stripe j on a Spanned VD	$S_j = \sum_{i=0}^{N-1} \{ \text{IF } J_i \geq j \\ \text{THEN } L_i \text{ ELSE } 0 \}$	$S_j = \sum_{i=0}^{N-1} \{ \text{IF } J_i \geq j \\ \text{THEN } L_i \text{ ELSE } 0 \}$
T_j	Number of blocks in stripes 0 through j	$T_j = \sum_{v=0}^j S_v$	$T_j = \sum_{v=0}^j S_v$
x	Offset of a data block from the beginning of a VD	0	$M-1$

Figure 35 gives an example of a Spanned SRL VD with five BVDs.

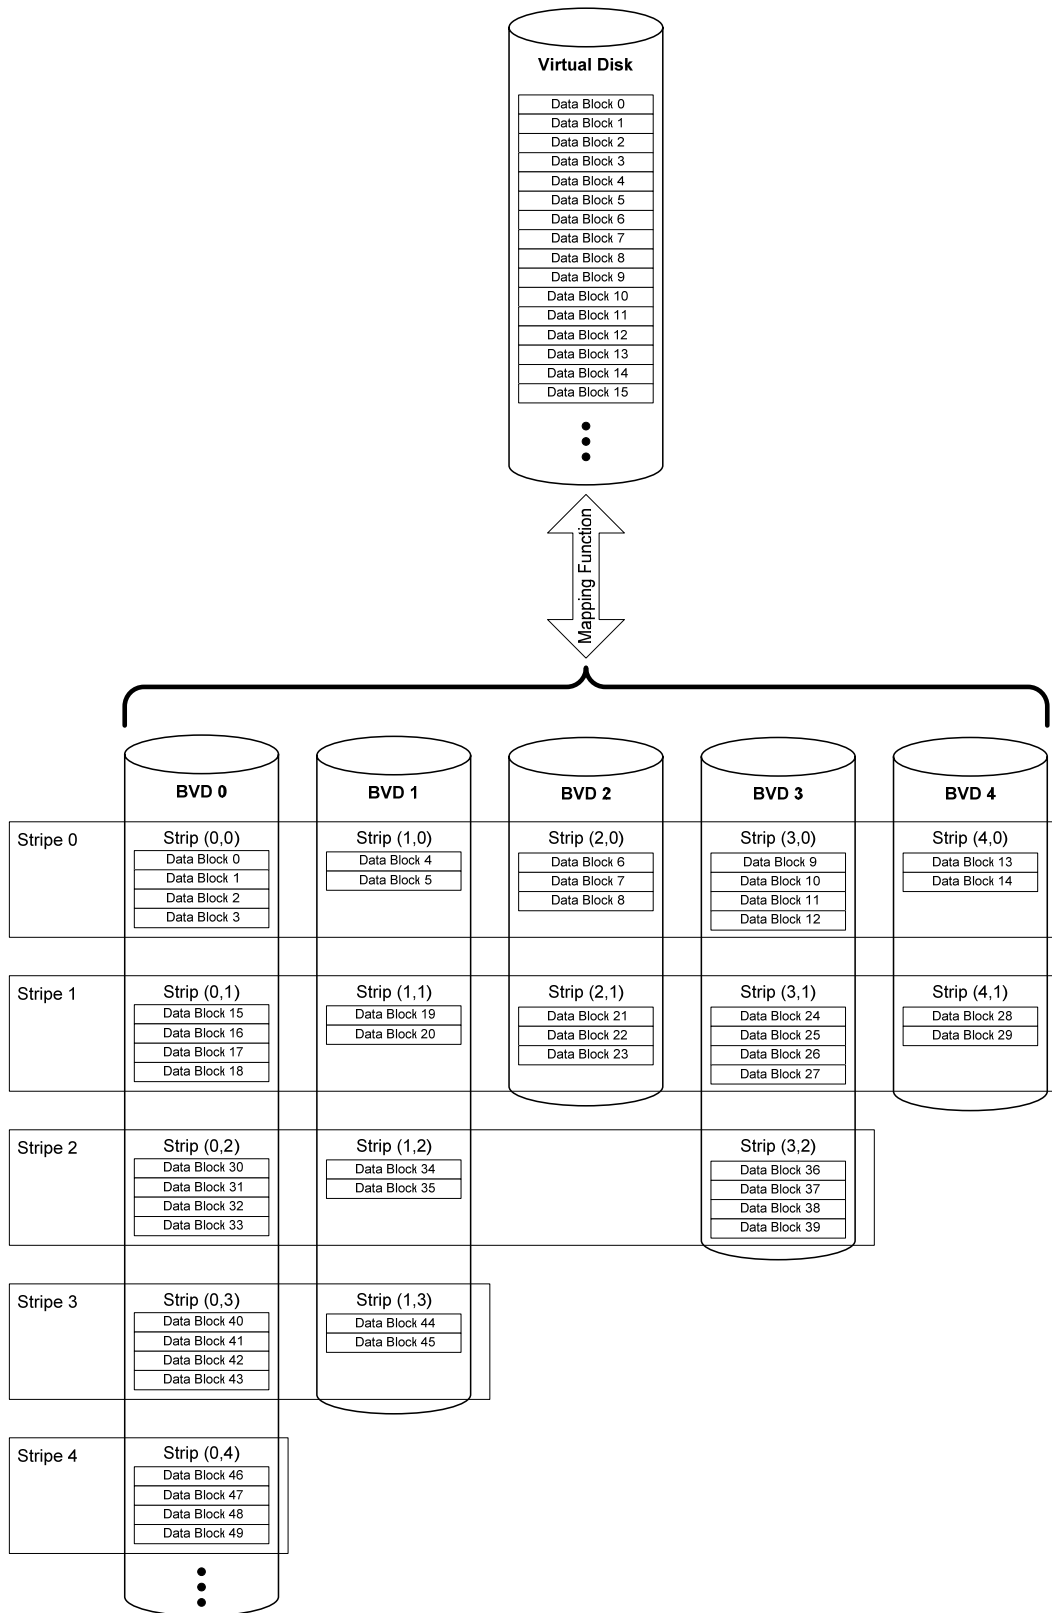


Figure 35: Spanned Secondary RAID Level (SRL=03)

5 DDF Structure

5.1 DDF Structure Overview

A DDF structure MUST reside on every physical disk that participates in a RAID configuration in a RAID storage subsystem. A minimum of 32MB MUST be reserved on each physical disk for a DDF structure. The last block of the reserved space MUST be the last addressable block of the physical disk.

Figure 36 illustrates the conceptual format of the DDF Structure on a physical disk with M addressable blocks. The DDF structure contains nine section types that are listed in Table 19.




Figure 36: DDF Structure

Table 19: DDF Sections

Contents	Context
DDF Header	Global
Controller Data	Global
Physical Disk Records	Global
Virtual Disk Records	Global
Configuration Records	Local
Physical Disk Data	Local
Bad Block Management Log	Local
Diagnostic Space	Local
Vendor Specific Logs	Vendor specific

The DDF Header section has three types. The Anchor DDF Header MUST be stored at the end of the physical disk. The last block of the Anchor DDF Header MUST be the last addressable logical block on a physical disk. The Anchor DDF Header contains a pointer to the Primary DDF Header. If DDF structure redundancy is implemented, the Anchor DDF Header also contains a pointer to the Secondary DDF Header.

Figure 36 shows an example of a DDF structure with redundant entries. Redundancy is OPTIONAL but if it is implemented the values of all fields in all section types except DDF Header sections MUST be equal in both the primary and secondary locations. Furthermore, the offset of the beginning of secondary sections from the start of the Secondary DDF Header MUST be equal to the offset of the beginning of the primary sections from the start of the Primary DDF Header. There MUST be only one DDF structure stored on a physical disk. The physical disk MUST not be partitioned using a partition table or another mechanism. The partitioning of capacity is handled by the DDF structure using the concept of Virtual Disks. The term DDF structure includes the Anchor DDF Header, the Primary DDF Header, the Secondary DDF Header, and all DDF Sections described in Table 19.

Signatures and CRCs are used to delineate the sections of the DDF structure. Sequence numbers are used to manage transitions during configuration and the change/update process. Timestamps are used to provide chronological hints. Configuration groups, Controllers, Physical Disks and Virtual Disks are uniquely identified by using “Globally Unique IDs” (GUIDs).


The size of sections like Physical Disk Records, Virtual Disk Records and Configuration Records can vary based on the maximum number of physical disks (PDs), virtual disks (VDs) and the maximum number of physical disks configurable within a virtual disk (PD per VD) a given implementation supports. This makes the DDF structure size variable.

A minimum of 32MB MUST be reserved on each physical disk to accommodate redundant DDF sections, optional spare blocks, and vendor specific logs or work space. It is NOT RECOMMENDED that redundant copies of DDF be placed adjacent to each other.

5.2 Byte Ordering

Each section of the DDF MUST be stored in big-endian format (i.e., the more significant bytes of the section are stored in lower addresses in relation to bytes of lesser significance). Figure 37 gives examples of the ordering of bytes in multi-byte fields.

Bit Ordering in a Byte



Byte Ordering in a 2-Byte Field



Byte Ordering in an M -Byte Field




Figure 37: Byte Ordering in Multi-Byte Fields

Multiple character ASCII strings are left justified when stored in multi-byte fields (i.e., the first character in the string is stored in the MSB of a multi-byte field). If an ASCII string has fewer characters than maximum allowed by the field in which it is to be stored, all unused bytes in the field **MUST** be filled with the ASCII character for “space.” Figure 38 gives the example of the 4-byte ASCII string “ABCD” stored in a 6-byte field.

4-Character ASCII String in a 6-Byte Field

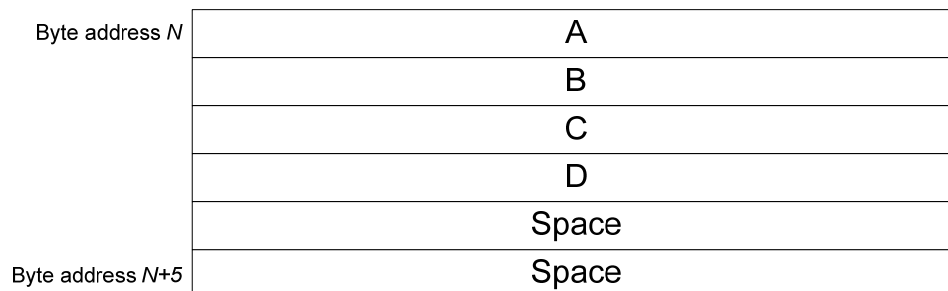


Figure 38: ASCII String Justification

Figure 39 gives an example of the byte ordering in the Controller Data Section of a DDF structure. The Controller Data Section is defined in Section 5.6 of this document.

Byte	Bit							
	7	6	5	4	3	2	1	0
0	(MSB)							
1	Controller Data Signature							
2								
3								
4								(LSB)
5	CRC							
6								
7								
8								(LSB)
9	Vendor ID							
10								
11								
12								
13								
14	Serial Number							
15								
16								
17								
18	Serial Number							
19								
20								
21								
22	Serial Number							
23								
24								
25								
26	Serial Number							
27								
28								
29								
30	Serial Number							
31								
32								
33								
34	PCI Vendor ID							
35								
36								
37								
38	PCI Device ID							
39								
40								
41								
42	PCI Sub Vendor ID							
43								
44								
45								
46	PCI Sub Device ID							
47								
48								
49								
50	Reserved							
51								
52								
53								
54	Vendor Unique Controller Data							
55								
56								
57								
58	Vendor Unique Controller Data							
59								
60								
61								
62	Vendor Unique Controller Data							
63								
64								
65								
Block_Size - 2	Vendor Unique Controller Data							
Block_Size - 1								(LSB)

Figure 39: Controller Data Byte Ordering Example

5.3 Signatures, Timestamps and CRCs

Signatures are unique 32-bit fields indicating the start of different sections in the DDF structure. Table 20 lists the signatures used by the DDF structure.

Table 20: DDF Signatures

Signature Name	Signature Value	Notes
DDF_Header	0xDE11DE11	
Controller_Data	0xAD111111	
Physical_Disk_Records	0x22222222	
Physical_Disk_Data	0x33333333	
Virtual_Disk_Records	0xDDDDDDDD	
VD_Configuration_Record	0xEEEEEEEE	
Spare_Assignment_Record	0x55555555	
VU_Configuration_Record	0x88888888	
Vendor_Specific_Log	0x01DBEEF0	
Bad_Block_Management_Log	0xABADB10C	

Timestamps are used to provide chronological hints for a number of DDF structures. A timestamp **MUST** be the number of seconds that have occurred since midnight GMT of January 1, 1980 and the time the timestamp is created.

The CRC **MUST** be calculated for a given structure with 0xFFFFFFFF as the initial and default value for the CRC field in the structure. The CRC **MUST** be calculated according to the CRC-32 algorithm described in ISO 3309 and in section 8.1.1.6.2 of ITU-T V.42. The CRC **MUST** be calculated using following polynomial:

$$X^{32}+X^{26}+X^{23}+X^{22}+X^{16}+X^{12}+X^{11}+X^{10}+X^8+X^7+X^5+X^4+X^2+X+1.$$

Once the CRC is calculated, the calculated value **MUST** replace the 0xFFFFFFFF stored in the CRC field during calculation.

5.4 GUIDs

The Globally Unique Identifier (GUID) is a structure used to uniquely identify the Controllers, Physical and Virtual Disks. It **MUST** be 24 bytes in length. A valid GUID **MUST NOT** use the values 0x20, 0x00 and 0xFF in Byte0 (MSB). These values in Byte 0 are reserved.

5.4.1 Controller GUID

The Controller GUID **MUST** be an ASCII string built by combining the T10 Vendor ID and the last 16 characters from the controller serial number. Information on T10 Vendor IDs and how to obtain one can be found at the T10 website (www.t10.org). Following is an example of a controller GUID, where the controller serial number is "11-10800BE10318" and the vendor name is "VendorID",

V	E	N	D	O	R	I	D		1	1	-	1	0	8	0	0	B	E	1	0	3	1	8
0							7	8															23

NOTE: If there is no serial number defined for the controller, then the 16 byte serial number field MUST be built by concatenating an 8-byte ASCII representation of a 4-byte timestamp (bytes 8-15) and an 8-byte ASCII representation of a 4-byte random number (bytes 16-23). This number SHOULD be created on the first configuration and stored in the controller's non-volatile memory. This field SHOULD NOT be user modifiable so that it remains constant for the controller's life from user's perspective.

The Physical Disk (PD) GUID MUST be 24 byte field.

[illegible]

H	D	D						2	0	0	4	0	2	0	1	A	A	B	B	C	C	D	D		
0								7		8														23	

Working Draft
Version 2.0 rev 19

following is an example of PD_GUID for an ATA disk where the serial number is “BB583GX3389103443379.”

A	T	A		B	B	5	8	3	G	X	3	3	8	9	1	0	3	4	4	3	3	7	9
0			3	4																			23

When a serial number is not available for ATA or SATA physical disks, the controller MUST generate one by concatenating an eight character current date in “yyyymmdd” ASCII format and a twelve byte hexadecimal ASCII representation of a six byte random number. The GUID generated using this serial number is considered a forced PD_GUID and the Force_PD_GUID_Flag in the PD_Type field of the Physical Disk Entry (Section 5.7.1) for the physical disk MUST be set. In addition, the Forced_PD_GUID_Flag in the disk’s Physical_Disk_Data section MUST also be set (Section 5.10). The following is an example of a forced PD GUID, where the date is February 1, 2004 and the twelve byte hexadecimal ASCII representation of the six byte random number is “AABBCCDDEEFF”

A	T	A		2	0	0	4	0	2	0	1	A	A	B	B	C	C	D	D	E	E	F	F
0			3	4																			23

5.4.3 Virtual Disk GUID

The VD GUID MUST be built by concatenating the creating controller’s T10 Vendor ID with a 16 byte unique identifier. The 16 byte identifier MAY be created by concatenating the Controller_Type field from Controller Data (Section 5.6), a 4-byte timestamp, and a 4-byte random number. Using this method, the VD_GUID provides data about age and the creating controller type. The following is an example of a VD_GUID with a vendor ID of “VENDORID”, a Controller_Type field of “AAAAAAA”, a time stamp of “BBBB”, and a random number of “CCCC.”

V	E	N	D	O	R	I	D	A	A	A	A	A	A	A	B	B	B	B	C	C	C	C
0							7	8							15	16			19	20		23

Alternatively, the unique identifier MAY be created by the controller using a vendor specific method. If this method is used, the controller MUST guarantee that the identifier is unique for all virtual disks in the system.

5.4.4 DDF Header GUID

The DDF Header GUID MUST be built by concatenating the creating controller’s T10 Vendor ID, the Controller_Type field from Controller Data (Section 5.6), a 4-byte timestamp, and a 4-byte random number. The DDF_GUID provides data about age and the creating controller type. The following is an example of a DDF_GUID with a vendor ID of “VENDORID”, a Controller_Type field of “AAAAAAA”, a time stamp of “BBBB”, and a random number of “CCCC.”

V	E	N	D	O	R	I	D	A	A	A	A	A	A	A	B	B	B	B	C	C	C	C
0							7	8							15	16			19	20		23

5.5 DDF Header

Three types of DDF header are defined: Anchor, Primary, and Secondary. The fields used in the DDF Headers are defined in Table 21. The Anchor, Primary and Secondary headers in the DDF structure

MUST have the same value in all fields except the CRC, Sequence_Number, Open_Flag and Header_Type. The Anchor header is NOT REQUIRED to track configuration changes and does not use the Sequence_Number and Open_Flag fields. All DDF Headers MUST adhere to the field definitions and uses described in Table 21.

DDF headers are global in context. Disks with different block sizes that participate in the same DDF structure will have differing values in the DDF Header fields whose values are based on a calculation related to block size.

Table 21: DDF Header Fields

Element	Size (Bytes)	Description
Signature	4	Unique Signature for DDF_Header Section (See Table 20)
CRC	4	CRC for the section (See Section 5.3). The CRC is covers the entire DDF Header section.
DDF_Header_GUID	24	GUID for the configuration or group (See Section 5.4.4)
DDF_rev	8	<p>This field contains the revision number of the DDF specification to which this DDF structure adheres. If the DDF structure adheres to an official version of the DDF specification, this field MUST contain an ASCII string in xx.yy.zz format, where x, y, and z are integer values 0-9. xx is the major revision number. yy is the minor revision number and zz is the development revision number. Minor DDF revisions are backward compatible to other minor revisions within the same major revision category.</p> <p>This field MAY also be used to state that the structure does not adhere to an official version of the DDF specification but is a vendor unique implementation of the structure. If the DDF structure uses a vendor unique specification, this field MUST be an ASCII string in CCxxxxxx format, where C is any non-numeric character. x can be any character.</p> <p>NOTE: The specification allows vendor unique implementations to be marked as such to allow a controller to notify the system administrator before importing or erasing configurations created using a vendor unique implementation of a DDF structure.</p> <p>Note: The DDF_rev for this version of the DDF specification is 02.00.00.</p>

Sequence_Number	4	Sequence number for configuration changes. Initial value MUST be 0x00000001. Wraparound is indicated by 0x00000000. Updating the Sequence_Number is OPTIONAL when the Vendor Specific Logs section, the Bad Block Management section or the Diagnostic Space section is updated. MUST be set to 0xFFFFFFFF in the Anchor header to indicate that configuration changes are not tracked in the Anchor Header.
Timestamp	4	Header update timestamp. MUST be set when the DDF header is updated. Used for a chronological hint only as this may not be consistent across controllers in the configuration.
Open_Flag	1	<p>DDF open/close flag. The DDF Open_Flag MUST be set on a physical disk before starting configuration change writes to the physical disk. The flag MUST be cleared (i.e., set to 0x00) after a successful write of configuration data to the physical disk. Open_Flag MUST NOT be set during log entry updates. This field MUST be set to 0xFF in the Anchor Header. The valid values for the flag are:</p> <p>0x00 = DDF closed; 0x01-0x0F = DDF Opened; 0x10-0xFE = Reserved; 0xFF = Not tracking DDF open/close. Used in the Anchor Header.</p> <p>NOTE: Any value of 0x01 through 0x0F means that the DDF structure is open. An implementation MAY use different values for different reasons the DDF structure is open. However, these reasons are not defined in the specification and may differ between different implementations.</p>

Foreign_Flag	1	<p>Foreign_Flag MAY be set when the controller has not imported the configuration defined by this DDF structure. With this flag set (and CRC recalculated/updated), nothing else is changed in the DDF structure and this physical disk MUST be marked as foreign until configuration is imported, cleared or the physical disk is physically removed. The valid values for this field are:</p> <p>0x00 = Not Foreign 0x01 = Foreign</p> <p>NOTE-1: This flag allows implementations where foreign configurations are only imported with a user's explicit acknowledgement. If a user has not acknowledged such an action then this flag is set and foreign configuration is not presented to OS level. This bit is persistent to tolerate controller swaps, denoting original foreign physical disks when the new controller sees all physical disks as foreign.</p> <p>NOTE-2: Some implementations prefer not to set the Foreign_Flag automatically. This allows a configuration to be moved back to the original controller without the original controller seeing the configuration as Foreign.</p>
Disk_Grouping	1	<p>Disk Grouping enforced/not enforced. When this field is set Disk Grouping MUST be enforced. The valid values for this field are:</p> <p>0x00 = Not Enforced 0x01 = Enforced</p>
Reserved	13	Filled with 0xFF
Header_Ext	32	Filled with 0xFF - Reserved for future use.
Primary_Header_LBA	8	LBA of the DDF Header of Primary DDF Structure.
Secondary_Header_LBA	8	LBA of the DDF Header of the Secondary DDF structure. This field MUST be filled with 0xFF if redundant DDF structures are not stored.
Header_Type	1	<p>0x00 = Anchor DDF Header; Stored on the last block of the physical disks 0x01 = Primary DDF Header 0x02 = Secondary DDF Header</p>
Reserved	3	Filled with 0xFF
Workspace_Length	4	Block count for reserved workspace that MAY be used for optional vendor specific functions. At least 16MB MUST be reserved for workspace. Thus, this field MUST have a value greater than or equal to 32,768/Block_Size where Block_Size is the block size of the physical disk on which this DDF_Header structure resides. Block size is stored in the Block_Size field of the Physical Disk Entry associated with the physical disk (See Section 5.7.1).

Workspace_LBA	8	LBA of the first block of the reserved workspace.
Max_PD_Entries	2	Maximum number of Physical Disk Entries. This is implementation dependent. Actual number of Physical Disks supported by a controller MAY be less than number of entries allowed. DDF structure implementations MUST only allow values of 15 (0x000F), 63 (0x003F), 255 (0x00FF), 1023 (0x03FF), and 4095 (0x0FFF).
Max_VD_Entries	2	Maximum number of Virtual Disk Entries. This is implementation dependent. Actual number of Virtual Disks supported by a controller MAY be less than number of entries allowed. DDF structure implementations MUST only allow values of 15 (0x000F), 63 (0x003F), 255 (0x00FF), 1023 (0x03FF), and 4095 (0x0FFF).
Max_Partitions	2	Maximum number of Configuration Record Entries (partitions) allowed per disk.
Configuration_Record_Length	2	<p>Virtual Disk Configuration Record (Section 5.9.1) length in blocks. Depends on the value of Max_Primary_Element_Entries field and calculated using the following formula:</p> $\text{ROUNDUP}((512 + (\text{Max_Primary_Element_Entries} * (4 + 8))) / \text{Block_Size})$ <p>Block_Size is the value in the Block_Size field of the Physical Disk Entry associated with the physical disk on which this DDF Header structure resides.</p>
Max_Primary_Element_Entries	2	<p>The maximum number of physical disks configurable within a basic VD. This field determines value of Configuration_Record_Length. DDF structure implementations MUST only allow values of 16 (0x0010), 64 (0x0040), 256 (0x0100), 1024 (0x0400), and 4096 (0x1000).</p> <p>NOTE: The value of this field determines value Configuration_Record_Length and size of two fields in the Virtual Disk Configuration Record: Physical_Disk_Sequence and Starting_Block.</p>
Max_Mapped_Block_Entries	4	The maximum number of supported Mapped_Block_Entries (See Section 5.11) in a physical disk's Bad Block Management Log. If the Bad Block Management Log section is not implemented, this field MUST be set to 0x00000000.
Reserved	50	Filled with 0xFF
Controller_Data_Section	4	Offset of the start of the Controller Data (Section 5.6) section from the Primary or Secondary DDF Header LBA
Controller_Data_Section_Length	4	Length of the Controller Data section in number of blocks
Physical_Disk_Records_Section	4	Offset of the start of the Physical Disk Records (Section 5.7) section from the Primary or Secondary Header LBA

Physical_Disk_Records_Section_Length	4	<p>Length of the Physical Disk Records section in number of blocks. This field is dependent on the entry in the Max_PD_Entries field. The value of this field is calculated according to the following formula:</p> $\text{ROUNDUP}(((\text{Max_PD_Entries} * 64) + 64) / \text{Block_Size})$ <p>Block_Size is the value in Block_Size field in the Physical Disk Entry associated with the physical disk on which this DDF Header structure resides.</p>
Virtual_Disk_Records_Section	4	Offset of the start of the Virtual Disk Records (Section 5.8) section from the Primary or Secondary DDF Header LBA
Virtual_Disk_Records_Section_Length	4	<p>Length of the Virtual Disk Records section in number of blocks. This value depends on the entry in the Max_VD_Entries field. The value of this field is calculated according to the following formula:</p> $\text{ROUNDUP}(((\text{MAX_VD_Entries} * 64) + 64) / \text{Block_Size})$ <p>Block_size is the value in the Block_Size field in the Physical Disk Entry associated with the physical disk on which this DDF Header structure resides.</p>
Configuration_Records_Section	4	Offset of the start of the Configuration Records (Section 5.9) section from the Primary or Secondary DDF Header LBA
Configuration_Records_Section_Length	4	<p>Length of the Configuration Records section in number of blocks. The value of this field MUST equal:</p> $\text{Configuration_Record_Length} * (\text{Max_Partitions} + 1).$
Physical_Disk_Data_Section	4	Offset of the start of the Physical Disk Data (Section 5.10) section from the start of the Primary or Secondary DDF Header LBA
Physical_Disk_Data_Section_Length	4	Length of the Physical Disk Data section in number of blocks
BBM_Log_Section	4	Offset of the start of the Bad Block Management Log (Section 5.11) from the start of the Primary or Secondary Header LBA. This is an OPTIONAL section and if the section is not implemented this field MUST be set to 0xFFFFFFFF.

BBM_Log_Section_Length	4	<p>Length of the Bad Block Management Log section in number of blocks. This is an OPTIONAL section and if the section is not implemented this field MUST be set to 0x00000000.</p> <p>If the BBM_Log_Section is implemented, the value of this field calculated according to the following formula:</p> $\text{ROUNDUP}(((\text{Max_Mapped_Block_Entries} * 16) + 32) / \text{Block_Size})$ <p>Block_Size is the value in Block_Size field in the Physical Disk Entry associated with the physical disk on which this DDF Header structure resides.</p>
Diagnostic_Space	4	Offset of the start of the Diagnostic Space (Section 5.11.1) section from the Primary or Secondary DDF Header LBA. This is an OPTIONAL section and if the section is not implemented this field MUST be set to 0xFFFFFFFF.
Diagnostic_Space_Length	4	Length of the Diagnostic Space section in number of blocks. This is an OPTIONAL section and if the section is not implemented this field MUST be set to 0x00000000.
Vendor_Specific_Logs_Section	4	Offset of the start of the Vendor Specific Logs (Section 5.13) from the Primary or Secondary DDF Header LBA. This is an OPTIONAL section and if the section is not implemented this field MUST be set to 0xFFFFFFFF.
Vendor_Specific_Logs_Section_Length	4	Length of the Vendor Specific Logs section in number of blocks. This is an OPTIONAL section and if the section is not implemented this field MUST be set to 0x00000000.
Reserved	256	Filled with 0xFF

5.6 Controller Data

This section is global in context and provides information about the last controller that operated on an attached RAID configuration. The Controller Data section MUST adhere to the field definitions and uses described in Table 22. When a new controller accesses a DDF structure previously operated on by another controller produced by the same vendor, it is up to the vendor implementation to determine whether or not the information stored in the Vendor_Unique_Controller_Data field is used by the new controller. When a new controller accesses a DDF Structure previously operated on by another controller produced by a different vendor, the new controller SHOULD erase the data contained in the Vendor_Unique_Controller_Data field and replace it with data specific to the new controller.

Table 22: Controller Data Fields

Element	Size (Bytes)	Description
Signature	4	Unique Signature for Controller_Data Section (See Table 20)
CRC	4	CRC for the section (See Section 5.3). The CRC covers the entire Controller_Data section.
Controller_GUID	24	Controller GUID (See Section 5.4.1)
Controller_Type	8	<p>For controllers with a PCI host interface, this field MUST contain the PCI ID of the controller.</p> <p>Byte 0-1: Vendor ID Byte 2-3: Device ID Byte 4-5: Sub Vendor ID Byte 6-7: Sub Device ID</p> <p>For controllers with a non-PCI host interface, bytes 6 and 7 of this field MUST be set to 0xFFFF. Bytes 0 through 5 MAY be set to any value at the discretion of the vendor.</p>
Product_ID	16	This field is an ASCII field. This field MUST contain the Product ID of the controller that last operated on the RAID configuration. For controllers that operate using the SCSI protocol, this field MUST be set to the value of the PRODUCT IDENTIFICATION field returned by the controller in the standard Inquiry data. For ATA controllers, this field MUST be set to the value the controller would return for the PRODUCT IDENTIFICATION field in the standard Inquiry data as described in the <u>SCSI / ATA Translation (SAT)</u> standard (INCITS 431-2007). For controllers that use neither the SCSI nor ATA protocols, the controller MUST create a Product_ID.
Reserved	8	Filled with 0xFF; Reserved for future use.
Vendor_Unique_Controller_Data	448	Vendor unique controller data.

5.7 Physical Disk Records

The Physical Disk Records section MUST list all configured physical disks attached to the controller. A physical disk is considered configured when it belongs to a Virtual Disk, is marked as a Spare, is configured as a Legacy (pass-through) disk or is a physical disk that needs to be tracked by the controller for a vendor-specific reason. This section is global in context (i.e., every configured physical disk has information about all configured physical disks participating in the system described by the DDF structure). This allows the detection of missing physical disks after a controller swap.

The length of this section is variable and is a function of maximum number of physical disks supported by the implementation. All Physical Disk Records sections MUST adhere to the field definitions and uses described in Table 23.

Table 23: Physical Disk Records Fields

Element	Size (Bytes)	Description
Signature	4	Unique Signature for Physical_Disk_Records Section (see Table 20).
CRC	4	CRC for the section (Section 5.3). The CRC covers the entire Physical_Disk_Records section.
Populated_PDEs	2	Number of Physical Disk Entries (PDE) used/populated
Max_PDE_Supported	2	Number of maximum PDE supported by vendor
Reserved	52	Filled with 0xFF
Physical Disk Entries	variable	64-byte PDEs. The number of PDE entries in this section is equal to Max_PD_Entries.

5.7.1 Physical Disk Entries

Physical Disk Entries (PDE) give minimal information about all physical disks attached to the controller. PDEs are stored in the Physical Disk Records section (Section 5.7) Each PDE is a 64 byte structure. Each PDE MUST adhere to the definitions and uses of the fields described in Table 24.

Table 24: Physical Disk Entry Fields

Element	Size (Bytes)	Description
PD_GUID	24	Physical Disk GUID (Section 5.4.2) for the physical disk. If this entry refers to the physical disk on which it is currently stored, the value of this field MUST contain the same value as the PD_GUID field in the Physical Disk Data section (Section 5.10). For unused entries, this field MUST be filled with 0xFF.
PD_Reference	4	Reference number to be used in VD Configuration Records (Section 5.9.1). This field MUST equal the value of the PD_Reference field of a physical disk's corresponding Physical Disk Data (Section 5.10)
PD_Type	2	Bit map describing a physical disk's type. Bit 0: 0 - Not using a forced PD_GUID (Section 5.4.2) 1 - Using a forced PD_GUID (also called the Forced_PD_GUID_Flag) Bit 1: 0 - Not participating in a VD 1 - Participating in a VD Bit 2: 0 - Not a global spare 1 - Global spare (VD Configuration Records are ignored) Bit 3: 0 - Not a spare disk 1 - Spare disk (Bit2 and Bit3 are exclusive. Bit3 MUST have precedence over Bit2) Bit 4: 0 - Not foreign 1 - Foreign (This is a Foreign disk and the Foreign_Flag in the DDF Header on this

		<p>disk may be set. See Section 5.5)</p> <p>Bit 5: 0 - Not a Pass-through/Legacy disk 1 - Pass-through/Legacy disk (No DDF structure is stored on this physical disk as DDF structures are only stored on configured physical disks and in the controller. If there are no other configured physical disks attached to the controller, then this information would only be stored in controller NVRAM. An implementation MAY restrict Pass-through/Legacy physical disks to systems with at least one configured disk attached to the RAID controller)</p> <p>Bit 6: Reserved Bit 7: Reserved Bits 8-11: Reserved Bits 15-12: Interface Type 0x0 = Unknown 0x1 = SCSI (parallel) 0x2 = SAS 0x3 = SATA 0x4 = FC 0x5-0xF = Reserved</p> <p>All reserved bits MUST be set to 0.</p>
PD_State	2	<p>State of the physical disk as part of one or more Virtual Disks.</p> <p>Bit 0: 0 - Offline 1 - Online Bit 1: 0 - OK 1 - Failed Bit 2: 0 - Not Rebuilding 1 - Rebuilding (Physical disk rebuilding for a failed physical disk) Bit 3: 0 - Not in Transition 1 - Transition (e.g., replacing a member physical disk through a copy operation) Bit 4: 0 - No PFA/SMART error alerts 1 - PFA/SMART error alerts Bit 5: 0 - No Un-recovered Read Errors 1 - Un-recovered Read Errors Bit 6: 0 - Not missing 1 - Missing (a physical disk may be Failed and Missing) Bit 7: Reserved Bits 8-15: Reserved</p> <p>All reserved bits MUST be set to 0.</p> <p>NOTE-1: Bit 1 MUST have precedence among Bits 0-3. If Bit 0 and Bit 1 are both set, then the physical disk's status is Failed and was probably Online before status change.</p> <p>NOTE-2: For a physical disk participating in</p>

		multiple VDs, the disk status Bits 1-3 MUST be set if any part of the physical disk is going through the corresponding process.
Configured_Size	8	Configured size of the physical disk in terms of highest addressable LBA during normal operation. The DDF structure MUST be stored at LBAs greater than or equal to the value in this field.
Path_Information	18	<p>This field contains information on the path from a controller to a physical disk or the attach point of the physical disk. This field is used for information purposes only.</p> <p>The Path_Information structure depends on the physical disk Interface Type as defined in Bits 12-15 of the PD_Type field.</p> <p>For Interface Type = 0x01 (SCSI), the Path_Information structure MUST be interpreted as follows:</p> <p>Bytes 0-3: Path 0 information</p> <ul style="list-style-type: none"> Bits 7-0: LUN Bits 15-8: SCSI Target ID Bits 23-16: SCSI Channel Bits 30-24: Reserved Bit 31: 0 - Path 0 OK 1 - Path 0 Broken <p>Bytes 4-7: Path 1 information (for systems with multiple paths for redundancy)</p> <ul style="list-style-type: none"> Bits 7-0: LUN Bits 15-8: SCSI Target ID Bits 23-16: SCSI Channel Bits 30-24: Reserved Bit 31: 0 - Path 1 OK 1 - Path 1 Broken <p>Bytes 8-17: Reserved</p> <p>For Interface Type = 0x2 (SAS), the Path_Information structure MUST be interpreted as follows:</p> <p>Bytes 0-7: Path 0 - SAS address of the end device (edge expander or controller) where the physical disk is attached.</p> <p>Bytes 8-15: Path 1 - SAS address of the end device (edge expander or controller) where the physical disk is attached.</p> <p>Byte 16:</p> <ul style="list-style-type: none"> Bit 0-6: Path 0 PHY identifier Bit 7: 0 - Path 0 OK 1 - Path 0 Broken <p>Byte 17:</p> <ul style="list-style-type: none"> Bit 0-6: Path 1 PHY identifier Bit 7: 0 - Path 1 OK

		<p>1 - Path 1 Broken</p> <p>For Interface Type = 0x3 (SATA), the Path_Information structure MUST be interpreted as follows:</p> <p>Bytes 0-7: Path 0 - Address of the end device in a SAS domain (edge expander or controller) where the physical disk is attached.</p> <p>Bytes 8-15: Path 1 - Address of the end device in a SAS domain (edge expander or controller) where the physical disk is attached.</p> <p>Byte 16:</p> <p>Bit 0-6: Path 0 PHY identifier</p> <p>Bit 7: 0 - Path 0 OK 1 - Path 0 Broken</p> <p>Byte 17:</p> <p>Bit 0-6: Path 1 PHY identifier</p> <p>Bit 7: 0 - Path 1 OK 1 - Path 1 Broken</p> <p>All reserved bytes or bits MUST be set to 0.</p> <p>NOTE-1: Path 1 information bytes MUST be set to 0xFF when dual path support is not implemented for this physical disk.</p> <p>NOTE-2: For FC physical disks, this field is undefined and all bytes MUST be set to 0x00.</p>
Block_Size	2	Logical block size of the disk in bytes. DDF implementations MUST only allow the following values in this field: 512, 1024, and 4096. Values other than 512, 1024 or 4096 are allowed but these values indicate vendor unique implementations.
Reserved	4	Filled with 0xFF

5.8 Virtual Disk Records

The Virtual Disk Records section lists all the configured VDs tracked by the DDF structure. It has global context and MUST be stored on all configured physical disks. The length of this section is variable and is a function of the maximum number of configurable virtual disks supported by the system implementation. The Virtual Disk Records section MUST adhere to the field definitions and uses described in Table 25.

Table 25: Virtual Disk Records Fields

Element	Size (Bytes)	Description
Signature	4	Unique Signature for Virtual_Disk_Records Section (See Table 20)
CRC	4	CRC for the section (See Section 5.3). The CRC covers the entire Virtual Disk Records section.
Populated_VDEs	2	Number of Virtual Disk Entries used/configured
Max_VDE_Supported	2	Maximum number of VDEs supported by the implementation

Reserved	52	Filled with 0xFF
Virtual_Disk_Entries	variable	64 Byte VDEs. The number of VDEs in this field is equal to Max_VD_Entries.

5.8.1 Virtual Disk Entries

Virtual Disk Entries (VDE) contain information about the configured VD's. VDEs are stored in the Virtual Disk Records (Section 5.8). Each VDE MUST adhere to the definitions and uses described in Table 26.

Table 26: Virtual Disk Entry Fields

Element	Size (Bytes)	Description
VD_GUID	24	Virtual Disk GUID (Section 5.4.3). For unused Virtual Disk Entries, this field MUST be filled with 0xFF.
VD_Number	2	The VD_Number MAY be used to map directly to the IDs (e.g., Logical Unit Numbers) presented to the OS as the result of an enumerate targets type command (e.g., SCSI Report LUNs command). For internal SCSI RAID controllers, the VD_Number MAY correspond directly to the LUN number reported to the OS. However, for external RAID controllers that have LUN masking and LUN mapping, the VD_Number cannot be used for this purpose and the RAID controller will have to implement its own method for maintaining this information. The VD_Number field needs to be managed during import MUST be consistent across reboots. Valid range is 0-0x7FFF. 0x8000-0xFFFF is reserved.
Reserved	2	Reserved Filled with 0xFF
VD_Type	4	<p>Bitmap describing the VD's type.</p> <p>Bit 0: 0 - Private 1 - Shared (Disk Grouping MUST be enforced when this bit is set)</p> <p>Bit 1: 0 - No Disk Grouping 1 - Disk Grouping Enforced</p> <p>Bit 2: 0 - VD_Name in ASCII format 1 - VD_Name in Unicode format</p> <p>Bit 3: 0 - Owner ID Not Valid 1 - Owner ID Valid</p> <p>Bits 15-4: Reserved</p> <p>Bits 31-16: Primary Controller GUID CRC. Used as a hint for the last controller that owned a VD in a clustered environment. The value stored in these bits MUST be the 16-bit CRC of the Controller GUID.</p> <p>All reserved bits MUST be set to 0.</p> <p>NOTE-1: Bits 16-31 MUST be ignored when Bit 0 is clear.</p> <p>NOTE-2: If Bit 0 is set, the value of Bit 1 MUST be ignored since Disk Grouping is automatically enforced.</p>

		<p>NOTE-3: If the Disk_Grouping field in the DDF Header is set to 0x01, Bit 1 MUST be ignored since Disk Grouping is automatically enforced for all disks.</p> <p>NOTE-4: If Bit 1 is set, disk grouping MUST be enforced for all disks participating in the virtual disk. This bit allows disk grouping to be enforced on some drives in a DDF configuration and not enforced on other drives.</p> <p>NOTE-5: If Bit 2 is set, the value stored in the VD_Name field MUST be compliant with Version 4.0 of the Unicode standard.</p> <p>NOTE-6: If Bit 3 is clear, bits 16-31 are invalid. If Bit 3 is set, bits 16-31 are valid.</p>
VD_State	1	<p>Bitmap describing the state of the virtual disk.</p> <p>Bits 2-0:</p> <ul style="list-style-type: none"> 0x0 - Optimal (The VD is operating and has experienced no failures of the disks that comprise VD.) 0x1 = Degraded (The VD has experienced at least one failure of the disks that comprise the VD. One more disk failure could result in the VD being placed in the "Failed" state indicating data loss has occurred.) 0x2 = Deleted (The VD has been marked as deleted by the system.) 0x3 = Missing (The VD has been marked as missing by the system.) 0x4 = Failed (The VD has experienced enough failures of the disks that comprise the VD for unrecoverable data loss to occur.) 0x5 = Partially Optimal (The VD has experienced one or more disk failures. The VD can experience at least one more disk failure before it is placed in the "Degraded" state.) 0x6 = Offline (The VD is offline for maintenance. Data on the VD is not available.) 0x7 = Reserved <p>Bit 3: 0 - Not Morphing 1 - Morphing (The VD is performing a morphing activity: RAID level migration, online capacity expansion, shrinking, defragmenting, stripe size migration, etc.)</p> <p>Bit 4: 0 - VD Consistent 1 - VD Not Consistent</p> <p>Bits 7-5: Reserved</p> <p>All reserved bits MUST be set to 0.</p> <p>NOTE-1: When the Morphing bit (Bit 3) is set, this</p>

		<p>VD cannot be imported to another controller from a different vendor as the new controller will not be able to determine the state or type of morphing activity. It is RECOMMENDED that no voluntary migrations of physical disks between controllers from different vendors be allowed when VDs associated with these physical disks are in a Morphing state.</p> <p>NOTE-2: The VD Not Consistent bit (Bit 4) MUST be set when a controller cannot guarantee a VD is consistent. The term consistent designates the state when all writes to a VD by client computer system, that are acknowledged as successfully completed by the controller, have been correctly written to the VD, including any redundancy information (e.g., parity). This bit SHOULD be cleared for a VD on a clean shutdown of the system. The bit MAY also be cleared during idle periods after a controller has determined that all writes have been consistently written to the VD. The bit SHOULD NOT be used as an actual cache synchronization flag.</p>
Init_State	1	<p>Bits 1-0: Initialization State 0x00 = Not Initialized 0x01 = Initialization in Progress 0x02 = Fully Initialized 0x03 = Reserved</p> <p>Bits 5-2: Reserved</p> <p>Bits 7-6: User Access Mode 0x00 = Read/Write 0x01 = Reserved 0x02 = Read Only 0x03 = Blocked (User reads and writes denied)</p> <p>NOTE: The Initialization in Progress state MAY be used to indicate that the VD is being initialized but is still available for read and/or write access. Some controllers provide this capability.</p>
Partially_Optimal_Drive_Failures_Remaining	1	<p>This field is only valid for BVDs and when VD_State = 'Partially Optimal.' This field indicates the number of remaining drive failures the VD may experience before leaving the 'Partially Optimal' state and entering the 'Degraded State'</p>
Reserved	13	Reserved Filled with 0xFF
VD_Name	16	<p>This field MAY be used to contain a 16 byte ASCII or Unicode string that is the name of the virtual disk. Bit 2 of the VD_Type field MUST be used to determine the Unicode or ASCII format of this field. This field MAY match the volume name used by some operating systems. If this field is not used, all bytes MUST be set to zero.</p>

5.9 Configuration Records

The Configuration Records section is local in context (i.e., relevant to physical disks where it is stored). It is variable in length and the length is stored in the Configuration_Record_Section_Length field in the DDF Header (Section 5.5). The Configuration_Record_Section_Length is implementation dependent and MUST be managed during import/merge operations between different controller implementations. A physical disk MAY have multiple partitions. Thus, a physical disk MAY be simultaneously listed in multiple configuration records. These configuration records MUST only be stored on the relevant (participating) physical disks. This section contains multiple records and these records MUST be either:

- a. No Virtual Disk Configuration Records;
- b. All Virtual Disk Configuration Records;
- c. One Spare Assignment Record;
- d. All Virtual Disk Configuration Records and one a Spare Assignment Record; or
- e. All Vendor Unique Configuration Records.

The Spare Assignment Record MUST only be stored on the relevant physical disk. The number of records stored on a physical disk MUST be one plus the value of the Max_Partitions field in the DDF header (Section 5.5). One extra record is provided to allow writing a modified entry before invalidating original. In the event that a Configuration Record section contains Max_Partitions Virtual Disk Configuration Records and One Spare Assignment Record, there will not be enough space to allow writing a modified entry before invalidating the original. It is up to the implementation to determine how to handle this situation. In certain situations, an implementation MAY convert a revertible spare to invalidate the Spare Assignment Record to provide the extra record for writing the new configuration. Another method would be to copy the Spare Assignment Record to Workspace temporarily and to copy it back after the new configuration record is written. To invalidate a Virtual Disk Configuration Record, a Spare Assignment Record, or a Vendor Unique Configuration Record, 0xFFFFFFFF MUST be stored in the Signature field of the record.

5.9.1 Virtual Disk Configuration Record

The size of VD configuration record MUST be the same as Configuration_Record_Length indicated in the DDF header (Section 5.5). Each VD configuration record MUST adhere to the definitions and uses described in Table 27.

Table 27: Virtual Disk Configuration Record Fields

Element	Size (Bytes)	Description
Signature	4	<p>For used entries, this field MUST be set to the Unique Signature for VD_Configuration_Record (Table 20).</p> <p>For unused entries, this field MUST be set to 0xFFFFFFFF.</p> <p>NOTE: Invalidating a used entry MAY be accomplished by only writing 0xFFFFFFFF to this field. This allows the possibility of un-deleting a configuration entry.</p>
CRC	4	CRC of the VD Configuration Record (Section 5.3)

VD_GUID	24	VD GUID (Section 5.4.3)
Timestamp	4	Configuration or reconfiguration timestamp. If this field is different than the timestamp in the VD's VD GUID, the configuration has changed since the VD was created. New timestamps MUST be greater than existing timestamps. Timestamps SHOULD only be used as chronological hints since they may not be consistent across controllers
Sequence_Number	4	Sequence number for configuration record changes. The initial value of this field MUST be 0x00000001. Subsequent configuration changes MUST increment this field by one. A wraparound of the sequence number MUST be indicated by the value 0x00000000.
Reserved	24	Filled with 0Xff
Primary_Element_Count	2	The number of physical disks used in a basic VD.
Strip_Size	1	<p>Stripe depth in $(2^n) \times \text{Block_Size}$ bytes format where n is the value of the field.</p> <p>Example (Block_Size = 512):</p> <p>n=0 - 512B stripe depth n=1 - 1KB stripe depth n=2 - 2KB stripe depth n=3 - 4KB stripe depth n=7 - 64KB stripe depth n=11 - 1MB stripe depth etc.</p> <p>Example (Block_Size = 4096):</p> <p>n=0 - 4KB stripe depth n=1 - 8KB stripe depth n=2 - 16KB stripe depth n=3 - 32KB stripe depth n=7 - 512KB stripe depth n=11 - 8MB stripe depth</p>
Primary_RAID_Level	1	RAID level of the BVD as defined in Table 1.
RAID_Level_Qualifier	1	The RAID level qualifier as defined in Table 2.
Secondary_Element_Count	1	The number of BVDs in a VD with a secondary RAID level as defined in Section 4.3 (e.g., RAID 50). For VDs without a secondary RAID level, this field MUST be set to 1.
Secondary_Element_Seq	1	Position of current basic VD in secondary VD. Valid only if secondary element count > 1.
Secondary_RAID_Level	1	The secondary RAID level as defined in Table 15. Valid only if Secondary_Element_Count > 1.
Block_Count	8	This field applies to the physical disk on which the configuration record is stored. The field states the size in blocks of the partition on the physical disk that is participating in the VD described by this configuration record.

VD_Size	8	The size of the user addressable space in the virtual disk. The size is stated in number of blocks.
Block_Size	2	Logical block size of the virtual disk in bytes. All physical disks participating in this virtual disk MUST have the same value in the Block_Size fields of their associated Physical Disk Entries (See Section 5.7.1).
Rotate parity count	1	Number of stripes before parity rotation for RAID 5R stored here as a power of 2 Value is n, n=0 to 31 0xFF means not used (or n=0)
Reserved	5	Filled with 0xFF
Associated_Spares	32	This field contains eight 4-byte entries for associated spare physical disks. Each used entry MUST contain the PD_Reference defined in the Physical Disk Entry (Section 5.7.1) for the associated spare physical disk. Unused entries MUST be set to 0xFFFFFFFF. Bytes 0-3: Spare Entry 0 Bytes 4-7: Spare Entry 1 Bytes 8-11: Spare Entry 2 Bytes 12-15: Spare Entry 3 Bytes 16-19: Spare Entry 4 Bytes 20-23: Spare Entry 5 Bytes 24-27: Spare Entry 6 Bytes 28-31: Spare Entry 7 NOTE: This field is used to detect missing dedicated spares as Spares Assignment Records are local.

Cache Policies & Parameters	8	<p>Cache policies for the VD. Cache policies, algorithms and parameters are implementation dependent. Therefore, bytes 1 through 7 are vendor specific. Byte 0 is a bit field where the bits are defined as:</p> <p>Bit0: 0-WriteThrough 1-Writeback</p> <p>Bit1: 0-Always (ignored if Bit0=0) 1-Adaptive (ignored if Bit0=0)</p> <p>Bit2: 0-No Read Ahead 1-Read Ahead</p> <p>Bit3: 0-Always (ignored if Bit2=0) 1-Adaptive (ignored if Bit2=0)</p> <p>Bit4: 0-No write caching if battery low or not present 1-Write caching allowed if battery low or not present</p> <p>Bit5: 0-No write caching allowed 1-Write caching allowed</p> <p>Bit6: 0-No read caching allowed 1-Read caching allowed</p> <p>Bit7: 0-No vendor specific caching algorithm 1-Vendor specific caching algorithm</p> <p>NOTE-1: Bits 4-6 are master enable/disable settings for cache parameters.</p> <p>NOTE-2: If bit7 is set, then bits 0-3 SHOULD left clear and ignored. Bits 4-6 SHOULD still be used.</p> <p>NOTE-3: During Vendor migration, if bit7 is found set, then the user SHOULD be notified and bit7 SHOULD be cleared. Bits 0-3 SHOULD be set to controller default and bits 4-6 SHOULD be preserved.</p>
BG_Rate	1	<p>This field is used to assign background task priorities to individual VDs. Examples of background tasks are: initialization, RAID level migration, expansion, etc. If the field is set to 0xFF, no background task rate has been set for this VD and the controller defaults SHOULD be used. If the field has a value of 0x00 through 0xFA, the VD SHOULD be assigned this relative priority for all background tasks. 0x00 is the lowest priority and 0xFA is the highest.</p> <p>NOTE: The actual weighting of the background task priorities is implementation dependent.</p>
Reserved	3	Filled with 0xFF

MDF Parity Disks	1	<p>Number of Parity disks in this configuration. Minimum value is 2 and Maximum value is 127.</p> <p>This field is reserved if the Primary RAID Level for this configuration is NOT 0x7.</p>
MDF Parity Generator Polynomial	2	<p>This field is used to specify the polynomial used to generate the Galois Field for Parity computation for this configuration.</p> <p>Valid values are 0x11dh, 0x169h, 0x1e7h, 0x12bh, 0x165h, 0x163h, 0x18dh, 0x12dh, 0x15fh, 0x1c3h, 0x1a9h, 0x187h, 0x14dh, 0x1cfh, 0x1f5h, and 0x171h.</p> <p>This field is reserved if the Primary RAID Level Qualifier for this configuration is NOT 0x7.</p>
Reserved	1	
MDF Constant Generation Method	1	<p>This field is used to specify the method used to generate constants that are used to compute parity.</p> <p>Value is n, and where n can be one of the following values: 0x00 - Vandermonde Method as defined in this specification; or, 0x05 - Vendor Defined. All other values are reserved.</p> <p>This field is reserved if the Primary RAID Level Qualifier for this configuration is NOT 0x7.</p>
Reserved	47	Filled with 0xFF
Reserved	192	Filled with 0xFF
V0	32	Filled with 0xFF - Reserved for future use
V1	32	Filled with 0xFF - Reserved for future use
V2	16	Filled with 0xFF - Reserved for future use
V3	16	Filled with 0xFF - Reserved for future use
Vendor Specific Scratch Space	32	Vendor unique information. All bytes in this field MUST be set to 0xFF when not used by the implementation.

Physical_Disk_Sequence	vari- able	<p>This field gives the sequence of physical disks in the BVD. This field contains multiple 4-byte entries. Each used entry MUST be a PD_Reference from a Physical_Disk_Entry (Section 5.7.1). Unused entries MUST be filled with 0xFFFFFFFF.</p> <p>Used entries provide the primary element sequence in ascending order. For the BVD, the number of used entries MUST equal the Primary_Element_Count in the Virtual_Disk_Configuration_Record for the BVD.</p> <p>If a physical disk that was part of this VD has been removed from the VD, reconfigured as part of another VD, and then returned to the system containing the original VD, the PD_Reference entry for the returned physical disk MUST be set to 0x00000000. This is to prevent the controller from erroneously reading the data from the replaced drive as part of the current VD.</p> <p>NOTE: This is a variable size field. Its size is determined by the following formula:</p> $\text{Max_Primary_Element_Entries} * 4$
Starting_Block	vari- able	<p>This field gives the starting LBAs for the partitions of the physical disks participating in the BVD. This field is variable in size and is comprised of 8-byte entries.</p> <p>Each entry contains an 8-byte LBA. Each entry corresponds to the physical drive with the PD_Reference stored in the same entry index in Physical_Disk_Sequence. Each entry MUST contain the starting LBA of its corresponding physical disk's partition that is part of the BVD.</p> <p>Unused entries MUST be filled with 0xFFFFFFFFFFFFFFFF.</p> <p>NOTE: This is a variable size field. Its size is determined by the following formula:</p> $\text{Max_Primary_Element_Entries} * 8$

5.9.2 Vendor Unique Configuration Record

The vendor unique configuration record allows proprietary implementations within DDF structures. Its size MUST be the same as the Configuration_Record_Length field of the DDF header (Section 5.5). This

structure is vendor specific except for the first 32-bytes which MUST adhere to the definitions and uses described in Table 28.

Table 28: Vendor Unique Configuration Record Mandatory Fields

Element	Size (Bytes)	Description
Signature	4	Unique Signature for VU_Configuration_Record (Table 20)
CRC	4	CRC of the Vendor_Unique_Configuration_Record (Section 5.3)
VD_GUID	24	VD GUID (5.4.3)

5.9.3 Spare Assignment Record

A Spare Assignment Record MUST be present on a physical disk only if the physical disk is a spare. The length of this record MUST be equal to the value in the Configuration_Record_Length field of the DDF Header (Section 5.5). Spare Assignment Records MUST adhere to the field definitions and uses described in Table 29.

Table 29: Spare Assignment Record Fields

Element	Size (Bytes)	Description
Signature	4	Unique Signature for Spare_Assignment_Record (Table 20)
CRC	4	CRC of the Spare_Assignment_Record (Section 5.3)
Timestamp	4	
Reserved	7	0xFFFFFFFFFFFFFFFF

Spare_Type	1	Bit0: 0 - Global 1 - Dedicated Bit1: 0 - Committable 1 - Revertible Bit2: 0 - Not Active 1 - Active Bit3: 0 - No Enclosure Affinity 1 - Enclosure Affinity Bits 7-4: Reserved The reserved bits MUST be set to 0. NOTE-1: Committable spares (Bit1 = 0) become permanent members of VDs after a rebuild. Revertible (Bit1 = 1) spares revert back to Spare status after the replacement of the original failed physical disk. An import, merge, or roaming action MAY require commitment of an active (failed-over) revertible spare. NOTE-2: An active spare is currently host user data on a VD for failed physical disk. NOTE-3: A spare with Enclosure Affinity MAY only be used as a spare for VDs that reside on disks in the same enclosure as the spare. Keeping track of which disks are associated with a particular enclosure is implementation dependent.
Populated_SAEs	2	Number of Spare Assignment Entries used
Max_SAE_Supported	2	Maximum number of SAE supported by vendor.
Reserved	8	Filled with 0xFF
Spare_Assignment_Entries	variable	Packed 32-byte spare assignment entries for spare physical disks (See Section 5.9.3.1). All bytes of unused entries MUST be set to 0xFF. NOTE: This is a variable size field which depends on Configuration_Record_Length. It is equal to Configuration_Record_Length minus 32 bytes.

5.9.3.1 Spare Assignment Entry

Spare Assignment Entries identify the VDs associated with a spare disk. The meaning and use of the Spare Assignment Entries is defined by Spare Type field of the Spare Assignment Record. Details on the use of Spare Assignment Entries are given in Section 5.9.3.2. Spare Assignment Entries MUST adhere to the field definitions and uses described in Table 30.

Table 30: Spare Assignment Entry Fields

Element	Size (Bytes)	Description
VD GUID	24	VD GUID of a VD to which this spare physical disk is assigned

Secondary_Element	2	VD Secondary Element number to which this spare physical disk is assigned. If the spare is not assigned to a particular secondary element, this field MUST be set to 0xFFFF.
Reserved	6	Filled with 0xFF

5.9.3.2 Spare Types and Spare Assignment Entry Details

5.9.3.2.1 Global Spares (Spare Type: Bit0 = 0)

For non-active global spares (Spare_Type: Bit2 = 0), all spare assignment entries in the Spare Assignment Record MUST be unused.

For active revertible global spares (Spare_Type: Bit1 = 1, Bit2 = 1), a Spare Assignment Entry MUST be present for each VD for which the spare is currently hosting user data. When a Spare Assignment Entry is added to a disk's Spare Assignment Record, a Virtual Disk Configuration Record MUST be added on the spare disk for the VD indicated by the VD GUID in the Spare Assignment Entry. All Virtual Disk Configuration Records associated with the VD on other disks MUST be updated to reflect that the spare is now part of the VD and has replaced the failed disk. Bit2 of the PD_State field in the Physical Disk Entry associated with the spare disk MUST be set to 1 (Rebuilding) while a VD's data is being rebuilt on the spare disk.

For an active revertible global spare, when an original failed disk for which the spare is hosting user data is replaced, Bit3 of the PD_State field in the Physical Disk Entry associated with the replacement disk MUST be set to 1 (Transition) while a VD's data is being copied to the replacement disk. Once the copy operation is completed, all Virtual Disk Configuration Records associated with the VD indicated by the VD GUID in the Spare Assignment Entry MUST be updated to reflect that the spare is no longer part of the VD and that the replacement disk is now part of the VD. The Spare Assignment Entry associated with the VD GUID MUST be deleted. The Virtual Disk Configuration record on the active revertible global spare disk associated with the VD GUID MUST be deleted.

For committable global spares (Spare Type: Bit1 = 0), once the disk joins a VD, it becomes a permanent member of the VD and no Spare Assignment Record will be present for the VD.

5.9.3.2.2 Dedicated Spares (Spare Type: Bit0 = 1)

For dedicated spares, a Spare Assignment Entry MUST be present for each VD for which the dedicated spare may serve as a replacement disk.

For active dedicated spares (Spare_Type: Bit2 = 1), a Virtual Disk Configuration Record MUST be present for each VD for which the spare is replacing a failed disk. When a Virtual Disk Configuration Record is placed on an active dedicated spare, all Virtual Disk Configuration Records associated with the same VD on other disks MUST be updated to reflect that the spare is now part of the VD and has replaced the failed disk. Bit2 of the PD_State field in the Physical Disk Entry associated with the spare disk MUST be set to 1 (Rebuilding) while a VD's data is being rebuilt on the spare disk.

For an active revertible dedicated spare, when an original failed disk for which the spare is hosting user data is replaced, Bit3 of the PD_State field in the Physical Disk Entry associated with the replacement disk MUST be set to 1 (Transition) while a VD's data is being copied to the replacement disk. Once the copy operation is completed, all Virtual Disk Configuration Records associated with the VD indicated by the VD GUID in the Spare Assignment Entry MUST be updated to reflect that the spare is no longer part

of the VD and that the replacement disk is now part of the VD. The Virtual Disk Configuration record on the active revertible dedicated spare disk associated with the VD GUID MUST be deleted.

5.10 For committable dedicated spares (Spare Type: Bit1 = 0), once the disk joins a VD, it becomes a permanent member of the VD and the Spare Assignment Record on the spare associated with the VD MUST be deleted.

Physical Disk Data

This section is local in context (i.e., only relevant to the physical disk on which it is stored). It stores the Physical Disk GUID and the Reference Number that is used in the Physical Disk Entry (see Section 5.7.1) for this physical disk. If the Physical Disk GUID is not available (possibly due to a missing or inaccessible disk serial number) or not usable, then a GUID MUST be generated (see Section 5.4.2) and the Forced_PD_GUID_Flag set to 1. Physical Disk Data sections MUST adhere to the field definitions and uses described in Table 31.

Table 31: Physical Disk Data Fields

Element	Size (Bytes)	Description
Signature	4	Unique signature for Physical_Disk_Data (Table 20)
CRC	4	CRC of the Physical_Disk_Data section (Section 5.3)
PD_GUID	24	Physical Disk GUID for this physical disk (Section 5.4.2)
PD_Reference	4	<p>Physical Disk Reference Number.</p> <p>For disks accessed by SCSI commands, the reference number MUST be built by computing a 32-bit CRC of the data contained in EVPD page 80.</p> <p>For ATA and SATA disks, the reference number MUST be built by computing a 32-bit CRC of the data returned by the Identify Device command.</p> <p>The reference number is global across the controller configuration and each disk on the controller MUST have a unique reference number. In case of conflict, a new value MUST be calculated. The method is implementation dependent but MUST insure that all disks in the configuration have unique PD_Reference values. This field MUST be managed when importing configured disk groups to insure the PD_Reference values remain unique.</p> <p>The values 0x00000000 and 0xFFFFFFFF are reserved and MUST not be used as PD_Reference values.</p>
Forced_Ref_Flag	1	<p>0x00 = No 0x01 = Yes</p> <p>The Forced_Ref_Flag is set when the PD_Reference is generated in a manner different than described in the PD_Reference field description. This situation occurs in the event two or more PD_References for different disks have the same</p>

		value.
Forced_PD_GUID_Flag	1	0x00 = No 0x01 = Yes Once a Forced PD GUID is created for a physical disk (see Section 5.4.2), the Forced PD GUID MUST be associated with the physical disk as long as the disk participates in the configuration.
Vendor Specific Scratch Space	32	Vendor unique information. All bytes in this field MUST be set to 0xFF when not used by the implementation.
Reserved	442	Filled with 0xFF

5.11 Bad Block Management Log

This section is local in context (i.e., only relevant to the physical disk where it is stored). It is an OPTIONAL section that allows block mapping and reassignment for defective blocks or the marking of blocks that cannot be recovered during a rebuild or recovery operation due to a media error. The Bad Block Management (BBM) Log section must adhere to the field definitions and uses described in Table 32.

Table 32: Bad Block Management Log Fields

Element	Size (Bytes)	Description
Signature	4	Unique Signature for Bad_Block_Management_Log (Table 20)
CRC	4	CRC for the BBM Log (Section 5.3). The CRC covers the entire Bad_Block_Management_Log section.
Entry_Count	4	Count for valid mapped/marked block entries. When there are no mapped/marked blocks, this field MUST be set to 0x00000000. The maximum entry for this field is the smaller of Max_Mapped_Block_Entries or Reserved_Spare_Block_Count.
Reserved_Spare_Block_Count	4	Up to 0xFFFFFFFF. The value of this field is implementation dependent.
Reserved	8	Filled with 0xFF
First_Spare_LBA	8	LBA of first spare block.
Mapped_Block_Entries	variable	Packed 16-byte entries for mapped LBAs (see Section 5.11.1). The entries MUST be sequentially populated. Unused entries must be filled with 0xFF. Note: This is a variable sized field. Its size is determined by the following formula: $\text{Max_Mapped_Block_Entries} * 16$

5.11.1 Mapped/Marked Block Entry

A Mapped/Marked Block Entry identifies a contiguous group of blocks that are remapped to another section on the physical disk or marked as unrecoverable due to a media error. Mapped/Marked Block Entries MUST adhere to the field definitions and uses described in Table 33.

Table 33: Mapped/Marked Block Entry Fields

Element	Size (Bytes)	Description
Defective_Block_Start	8	The LBA of the first defective block in a contiguous group of defective blocks that are remapped or marked as unrecoverable.
Spare_Block_Offset	4	If this field is set to 0xFFFFFFFF, the blocks indicated by this Mapped/Marked Block Entry are marked as unreadable due to media error. If this field is not set to 0xFFFFFFFF, the blocks MUST be remapped to a new location. The field contains the offset from the First_Spare_LBA where the group of defect blocks MUST be remapped.
Remapped_Marked_Count	2	Number of blocks to be remapped or marked as unrecoverable.
Reserved	2	Filled with 0xFF

5.12 Diagnostic Space

The Diagnostic Space section is local in context. This section is OPTIONAL. If implemented, the size of this section is stated in the Diagnostic_Space_Length field of the DDF Header (Section 5.5). The section MAY be used for destructive read/write tests on the physical disk.

5.13 Vendor Specific Logs

The Vendor Specific Logs section is OPTIONAL. The structure of this section is implementation/vendor specific except for the first 32 bytes which are used for mandatory fields. The first 32 bytes MUST adhere to the field definitions and uses described in Table 34.

Table 34: Vendor Specific Logs Mandatory Fields

Element	(Size Bytes)	Description
Signature	4	Unique signature for Vendor_Specific_Logs section (Table 20)
CRC	4	CRC for the Vendor Specific Logs section (Section 0) The CRC is covers entire Vendor_Specific_Logs section. Any unused bytes in the section MUST be filled with 0xFF.
Log_Owner	8	T10 vendor ID of the log owner. If the vendor specific logs are used, this field MUST be populated with the T10 Vendor ID of the vendor that created the log. It MUST be set to 0xFFFFFFFFFFFFFFFF when the logs are empty or not used.
Reserved	16	Filled with 0xFF

