To: X3T10 Committee (SCSI)

From: George Penokie (IBM)

Subject: More Simplification of SACL Configuration

1 SCSI Implementation

This implementation of the simplified configuration of a storage subsystem will use the CREATE/MODIFY STORAGY ARRAY CONFIGURATION service action to deliver the configuration parameters to the storage subsystem.

1.1 CREATE/MODIFY STORAGE ARRAY CONFIGURATION service action

The CREATE/MODIFY STORAGE ARRAY CONFIGURATION service action (see table 1) requests the creation of a new volume set and redundancy group, or the modification of an existing volume set and redundancy group. If the create operation fails to complete successfully the command shall be terminated with a CHECK CONDITION status. The sense key shall be set to HARDWARE ERROR, and the additional sense code set to CREATION OF LOGICAL UNIT FAILED. If the modification operation fails to complete successfully the command shall be terminated with a CHECK CONDITION status. The sense key shall be set to HARDWARE ERROR, and the additional sense code set to MODIFICATION OF LOGICAL UNIT FAILED.

Table 1 - CREATE/MODIFY STORAGE ARRAY CONFIGURATION service action

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (BFh)							
1	RESERVED			SERVICE ACTION (08h)				
2	REDUNDANCY TYPE IDENTIFIER							
3	BUSPROC	RESERVED		EQSPRD	RESERVED			
4	(MSB)	LUN_V						
5								(LSB)
6	(MSB)	LIST LENGTH						
7								
8								
9								(LSB)
10	CREATE/MODIFY		CONFIGURE		RESERVED			IMMED
11	CONTROL							

The REDUNDANCY TYPE IDENTIFIER field indicates the type of protection that shall be used within the redundancy group being created or modified. See table 2 for the format of the REDUNDANCY TYPE IDENTIFIER field.

Table 2 - REDUNDANCY TYPE IDENTIFIERS

Codes	Description
00h	No redundancy
01h	Copy redundancy
02h	XOR redundancy
03h	P+Q redundancy
04h	P+S redundancy
05h	S redundancy
06h-7Fh	Reserved
80h-FFh	Vendor specific

The redundancy type identifier field shall only indicate the minimum amount of protection required by the redundancy group being configured or modified. See table 3 for the minimum requirements for each type of redundancy.

Table 3 - Minimum redundancy group protection

Type	Minimum protection within the configured redundancy group
No redundancy	The SACL is not required to protect user data.
Copy redundancy	The SACL shall duplicate all user data at least one time, preferably on different peripheral devices.
XOR redundancy	The SACL shall protect user data such that a single peripheral device failure does not cause loss of user data.
P+Q redundancy	The SACL shall protect user data such that the failure of two peripheral devices does not cause loss of user data.
P+S redundancy	The SACL shall protect user data such that the failure of a single peripheral device does not cause loss of user data and after some vendor specific amount of time a second peripheral device failure shall not cause loss of user data. If a second peripheral device fails within the vendor specific amount of time user data may be lost.
S redundancy	The SACL shall have access to at least enough spare space to allow a rebuild of one peripheral device within that space. User data may be lost if the SACL cannot predict or predicts incorrectly a peripheral device failure. User data may also be lost if the SACL does not complete a rebuild before the peripheral device fails.

An equal user data spreading (EQSPRD) bit of zero indicates the target may spread user data in a nonuniform manner over the peripheral devices associated with the volume set being created or modified. A EQSPRD bit of one indicates the target shall spread user data in a uniform manner over all the peripheral devices associated with the volume set being created or modified.

A bus protection (BUSPROC) bit of zero indicates that the target be configured such that a single bus failure may cause the application client to lose access to user data within the volume set being created or modified. A BUSPROC bit of one indicates that the target shall be configured so a single bus failure shall not cause the application client to lose access to any user data within the volume set being created or modified.

The LUN_V field specifies the address of the volume set that shall be created or modified.

An immediate (IMMED) bit of zero indicates that status shall be returned after the create/modify storage

array operation has completed. An IMMED bit of one indicates that the storage array shall return status as soon as the command descriptor block has been validated, and the entire CREATE/MODIFIY STORAGE ARRAY CONFIGURATION parameters list has been transferred.

The CONFIGURE field is defined in table 4.

Table 4 - CONFIGURE

Codes	Description
00b	Any unassigned p_extent(s) within the target that received the CREATE/MODIFY STORAGE ARRAY CONFIGURATION service action may be used to configure the selected volume set and redundancy group to the requested capacity. Any CREATE/MODIFY PERIPHERAL DEVICE DESCRIPTORS (table 10) shall be ignored.
01b	The target shall use the CREATE/MODIFY STORAGE ARRAY CONFIGURATION parameter list (table 6) to determine the configuration of the volume set and redundancy group. The EQSPRD bit shall be ignored.
10b	All unassigned p_extents within the target that received the CREATE/MODIFIY STORAGE ARRAY CONFIGURATION service action shall be configured into a volume set and a redundancy group. The VOLUME SET CAPACITY field (table 6) and any CREATE/MODIFY PERIPHERAL DEVICE DESCRIPTORS (table 10) shall be ignored.
11b	Reserved

The CREATE/MODIFY field is defined in table 5.

Table 5 - CREATE/MODIFY

Codes	Description
00b	The target shall create a volume set and a redundancy group and shall assign to the created volume set the logical unit number contained in the LUN_V field. The target shall assign to the created redundancy group a logical unit number per the addressing rules (xxx). If the addressed volume set already exists within the target the target shall modify the existing volume set and its associated redundancy group as requested in the CREATE/MODIFY STORAGE ARRAY CONFIGURATION service action. The target may preserve the contents of and access to user data on completion of a modify.
01b	The target shall create a volume set and a redundancy group, and shall assign to the created volume set and redundancy group logical unit numbers per the addressing rules (xxx). The LUN_V field shall be ignored.
10b	The target shall modify the volume set addressed in the LUN_V field and its associated redundancy group. If the addressed volume set does not exist the target shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code set to LOGICAL UNIT NOT CONFIGURED.The target may preserve the contents of and access to user data on completion of the modify.
11b	The target shall modify the volume set addressed in the LUN_V field and its associated redundancy group. If the addressed volume set does not exist the target shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code set to LOGICAL UNIT NOT CONFIGURED. The target shall preserve the contents of and access to user data on completion of the modify.

The CREATE/MODIFY STORAGE ARRAY CONFIGURATION parameter list (see table 6) contains user data mapping information and a list of CREATE/MODIFY PS_EXTENT DESCRIPTORS that are used to create or modify the addressed volume set and its associated redundancy group.

Table 6 - CREATE/MODIFY STORAGE ARRAY CONFIGURATION parameter list

Bit / Byte	7	6	5	4	3	2	1	0
0	(MSB)							
1				CAPACITY				
2								
3								(LSB)
4	(MSB)							
5				BYTES PER BLOCK				(LSB)
6	(MSB)							
7				NORMAL USER DATA TRANSFER SIZE				(LSB)
8	RESERVED							
9	REBUILD/RECALCULATE PRIORITY							
10	PERCENTAGE OF SEQUENTIAL READ TRANSFERS							
11	PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS							
	CREATE/MODIFY PERIPHERAL DEVICE DESCRIPTORS(S) (if any)							
12				CREATE/MODIFY PERIPHERAL DEVICE DESCRIPTOR 0				
15								
	. . .							
n-3				CREATE/MODIFY PERIPHERAL DEVICE DESCRIPTOR X				
n								

The CAPACITY field contains the size to configure the volume set and the redundance group in logical blocks. If the CREATE/MODIFY field is 10b the new size of the volume set being modified shall be set to the value in the CAPACITY field and the new size of the redundancy group shall be set to the value in the CAPACITY field.

NOTE 1 - Attempts by an application client to modify a volume set to a smaller capacity may result in a loss of user data.

The BYTES PER BLOCK field contains the size, in bytes, of the logical blocks in the CAPACITY field and the NORMAL USER DATA TRANSFER SIZE field. A value of zero in the BYTES PER BLOCK field shall indicate the number of bytes per logical block is 512.

The NORMAL USER DATA TRANSFER SIZE field contains the number of logical blocks the application client normally requests transferred on each user data transfer. The target shall treat the NORMAL USER DATA TRANSFER SIZE field as an advisory parameter. A NORMAL USER DATA TRANSFER SIZE field of zero indicates

the application client is providing no direction on the size of user data transfers.

The REBUILD/RECALCULATE PRIORITY field contains the length of time the target should take to do a rebuild operation or a recalculate operation. The target shall treat the REBUILD/RECALCULATE PRIORITY field as an advisory parameter. If the REDUNDANCY TYPE IDENTIFIER field contains a zero (i.e., no redundancy) then the target shall ignore the REBUILD/RECALCULATE PRIORITY field. See table 7 for a description of the contents of the REBUILD/RECALCULATE PRIORITY field.

Table 7 - Rebuild/recalculate priority selection

Codes	Description
00h	The application client is providing no direction on the length of time for rebuilds or recalculates
01h	The target shall suspend rebuild and recalculate operations during all read/write requests from any application client.
02h-EFh	An indication of the length of time the target should take to do a rebuild operation or a recalculate operation. Generally, larger values indicate shorter rebuild and recalculate times. NOTE 2 - The effect of different rebuild/recalculate times is to increase and decrease the performance of a target. Lower values increase performance but at a cost of being exposed to data loss for a longer time. Higher values decrease performance but keep the exposure to data loss at a minimum.
FFh	The target shall not accept any read/write requests from an application client until the rebuild or recalculate operation is complete. All attempts to read or write shall be l terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code set to LOGICAL UNIT NOT READY, REBUILD IN PROGRESS or LOGICAL UNIT NOT READY, RECALCULATION IN PROGRESS.

The PERCENTAGE OF SEQUENTIAL READ TRANSFERS field contains the percent of times an application client is expected to do reads of sequential logical blocks on consecutive user data read transfers. The target shall treat the PERCENTAGE OF SEQUENTIAL READ TRANSFERS field as an advisory parameter. See table 8 for a description of the contents of the PERCENTAGE OF SEQUENTIAL READ TRANSFERS field.

Table 8 - PERCENTAGE OF SEQUENTIAL READ TRANSFERS

Codes	Description
0	The application client is providing no direction on the sequentially of user data read transfers.
1-100	The percent of times an application client is expected to do reads of sequential logical blocks on consecutive user data read transfers.
101-127	Reserved
128-255	Vendor Specific

The PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS field contains the percent of times an application client is expected to do writes of sequential logical blocks on consecutive user data write transfers. The target shall treat the PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS field as an advisory parameter. See table 9 for a description of the contents of the PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS field.

TABLE 9 - PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS

Codes	Description
0	The application client is providing no direction on the sequentially of user data write transfers.
1-100	The percent of times an application client is expected to do writes of sequential logical blocks on consecutive user write read transfers.
101-127	Reserved
128-255	Vendor Specific

The CREATE/MODIFY PERIPHERAL DEVICE DESCRIPTOR contains information the target shall use to control the user data mapping within peripheral devices. See table 10 for the format of the CREATE/MODIFY PERIPHERAL DEVICE DESCRIPTOR.

Table 10 - Data format of CREATE/MODIFY PERIPHERAL DEVICE DESCRIPTOR

Bit Byte	7	6	5	4	3	2	1	0
0	LUN_P							
1								
2	RESERVED							
3	WEIGHTING OF USER DATA							

The LUN_P field defines the address of the peripheral device to place user data.

The WEIGHTING OF USER DATA field contains a value used to calculate the portion of the volume set's and redundancy group's capacity to place on the selected peripheral device. The target shall determine this capacity by using the equation $c = (dw) \times [(vc)/(\Sigma dw)]$ where:

c = capacity of the peripheral device selected in the LUN_P field,
dw = value of WEIGHTING OF USER DATA field, and
vc = value of the CAPACITY field.

If the requested capacity (c) will not fit within the unassigned p_extent area(s) on the addressed peripheral device the target shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code set to PARAMETER VALUE INVALID.

1.2 REPORT STORAGE ARRAY CONFIGURATION service action

The REPORT STORAGE ARRAY CONFIGURATION service action (see table 11) requests that information regarding the SCSI storage array's configuration be sent to the application client. If this service action requests information on a volume set that has more than one associated redundancy group the target shall terminate the command with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code set to INVALID FIELD IN CDB.

Table 11 - REPORT STORAGE ARRAY CONFIGURATION service action

Bit Byte	7	6	5	4	3	2	1	0
0	OPERATION CODE (BEh)							
1	RESERVED			SERVICE ACTION (02h)				
2	RESERVED							
3	RESERVED							
4	(MSB)			LUN_V				
5								(LSB)
6	(MSB)							
7				ALLOCATION LENGTH				
8								
9								(LSB)
10	RESERVED							
11	CONTROL							

The LUN_V field specifies the address of the volume set for which information shall be reported per table 12. If the requested logical unit has not been configured the command shall be terminated with a CHECK CONDITION status. The sense key shall be set to ILLEGAL REQUEST, and the additional sense code set to LOGICAL UNIT NOT CONFIGURED.

The REPORT CONFIGURATION parameter list is defined in table 12.

Table 12 - REPORT CONFIGURATION parameter list

Bit Byte	7	6	5	4	3	2	1	0
0	RESERVED							
1	REDUNDANCY TYPE IDENTIFIER							
2	BUSPROC	RESERVED		EQSPRD	RESERVED			
3	RESERVED	STATE OF THE VOLUME SET						
4	(MSB)	CAPACITY						
5								
6								
7								(LSB)
8	(MSB)	BYTES PER BLOCK						
9								(LSB)
10	(MSB)	NORMAL USER DATA TRANSFER SIZE						
11								(LSB)
12	RESERVED							
13	REBUILD/RECALCULATE PRIORITY							
14	PERCENTAGE OF SEQUENTIAL READ TRANSFERS							
15	PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS							
16	RESERVED							
17	RESERVED							
18	(MSB)	VOLUME SET PERIPHERAL DEVICE DESCRIPTOR LIST LENGTH						
19								(LSB)
	VOLUME SET PERIPHERAL DEVICE DESCRIPTOR(S)							
20		VOLUME SET PERIPHERAL DEVICE DESCRIPTOR (First)						
23								
	. . .							
n-3		VOLUME SET PERIPHERAL DEVICE DESCRIPTOR (Last)						
n								

The REDUNDANCY TYPE IDENTIFIER field (table 2) indicates the type of protection being used within the redundancy group associated with the addressed volume set. For a description of the redundancy group methods see xxx.

An equal user data spreading (EQSPRD) bit of zero indicates the target is configured such that the user data is spread in a nonuniform manner over the peripheral devices associated within the addressed volume set. An EQSPRD bit of one indicates the target is configured such that the user data is spread in a uniform manner over all the peripheral devices associated within the addressed volume set.

A bus protection (BUSPROC) bit of zero indicates that the target is configured such that a single bus failure causes the application client to lose access to user data within the addressed volume set. A BUSPROC bit of one indicates that the target is configured such that single bus failure does not cause the application client to lose access to any user data within the addressed volume set.

The VOLUME SET STATE field is defined in xxx.

The CAPACITY field indicates the size of the addressed volume set in logical blocks.

The BYTES PER BLOCK field indicates the size, in bytes, of the logical blocks in the CAPACITY field and the NORMAL USER DATA TRANSFER SIZE field.

The NORMAL USER DATA TRANSFER SIZE field indicates the number of logical blocks the target expects during a normal user data transfer. A NORMAL USER DATA TRANSFER SIZE field of zero indicates the target has received no direction on the size of user data transfers.

The REBUILD/RECALCULATE PRIORITY field indicates the length of time the target takes to do a rebuild operation or a recalculate operation. See table 13 for a description of the contents of the REBUILD/ RECALCULATE PRIORITY field.

Table 13 - Rebuild/recalculate priority selection

Codes	Description
00h	Shall indicate the target has received no direction on how long it will take to do rebuilds nor recalculates or that the associated redundancy group is configured as no redundancy.
01h	Indicates the target shall suspend rebuild and recalculate operations during all read/write requests from any application client.
02h-EFh	An indication of the length of time the target takes to do a rebuild operation or a recalculate operation. Generally, larger values indicate shorter rebuild and recalculate times.
FFh	Indicates the target shall not accept any read/write requests from an application client until the rebuild or recalculate operation is complete.

The PERCENTAGE OF SEQUENTIAL READ TRANSFERS field indicates the percent of times the target expects reads of sequential logical blocks on consecutive user data read transfers. See table 14 for a description of

the contents of the PERCENTAGE OF SEQUENTIAL READ TRANSFERS field.

Table 14 - PERCENTAGE OF SEQUENTIAL READ TRANSFERS

Codes	Description
0	The target has received no direction on the sequentially of user data read transfers.
1-100	The percent of times the target expects reads of sequential logical blocks on consecutive user data read transfers.
101-127	Reserved
128-255	Vendor Specific

The PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS field indicates the percent of times the target expects writes of sequential logical blocks on consecutive user data write transfers. See table 15 for a description of the contents of the PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS field.

TABLE 15 - PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS

Codes	Description
0	The target has received no direction on the sequentially of user data write transfers.
1-100	The percent of times the target expects writes of sequential logical blocks on consecutive user write read transfers.
101-127	Reserved
128-255	Vendor Specific

The VOLUME SET PERIPHERAL DEVICE DESCRIPTOR contains a list of peripheral devices associated with the addressed volume set. See table 16 for the format of the VOLUME SET PERIPHERAL DEVICE DESCRIPTOR field.

Table 16 - VOLUME SET PERIPHERAL DEVICE DESCRIPTOR

Bit Byte	7	6	5	4	3	2	1	0
0	LUN_P							
1								
2	RESERVED							
3	WEIGHTING OF USER DATA							

The LUN_P field indicates the address of a peripheral device associated with the addressed volume set.

The WEIGHTING OF USER DATA field indicates the portion of the volume set's capacity and the redundancy group's placed on the selected peripheral device.

NOTE 3 - The WEIGHTING OF USER DATA fields may contain the same values as the target received from the application client or it may be calculated using an equation. One example of an equation that may be used by the target to calculate the value in the WEIGHTING OF USER DATA field is $dw = (vc)/c$ where:

c = capacity of the peripheral device selected in the LUN_P field assigned to the selected volume set,

dw = value of WEIGHTING OF USER DATA field, and
vc = value of the CAPACITY field.

Annex A Example of a SCSI storage array configuration using a CREATE/MODIFY ARRAY CONFIGURATION service action

This annex is only intended to be an example of how a storage array could use the parameters received in a CREATE/MODIFY STORAGE ARRAY CONFIGURATION service action. Each storage array will have different factors that will effect the way it uses these parameters to configure volume sets and redundancy groups. These factors would include things like:

a) the number of peripheral devices within the storage array,
b) the capacity of the peripheral devices,
c) the ratio of link speed vs. media speed,
d) the ratio of access time vs. link speed, etc.

On receipt of a CREATE/MODIFY STORAGE ARRAY CONFIGURATION service action the target will examine the parameters and configure a volume set and a redundancy group using those parameters.

The contents of the NORMAL USER DATA TRANSFER SIZE field and the PERCENTAGE OF SEQUENTIAL TRANSFERS field in combination are used by the target to determine the user data stripe depth mapping.

See table 17 for an example of user data stripe depths the target in this example selects based on the contents of the NORMAL USER DATA TRANSFER SIZE field and PERCENTAGE OF SEQUENTIAL READ TRANSFERS field. In this example the PERCENTAGE OF SEQUENTIAL WRITE TRANSFERS field is ignored.

Table 17 - User data stripe depth mapping selection

		PERCENTAGE OF SEQUENTIAL READ TRANSFERS			
		1% to 25%	25% to 50% or 0%	51% to 75%	greater than 75%
NORMAL USER DATA TRANSFER SIZE in logical blocks	1 to 128	1	2		3
	129 to 512 or 0	2		4	
	greater than 512	3	4		5

Notes:
1) user data stripe depth = 4 x normal user data transfer size
2) user data stripe depth = 2 x normal user data transfer size
3) user data stripe depth = normal user data transfer size
4) user data stripe depth = normal user data transfer size / (number of disk drives/2)
5) user data stripe depth = normal user data transfer size / number of disk drives

The contents of the REBUILD/RECALCULATE PRIORITY field is used by the target to determine the how long

the a rebuild or recalculate operation will take to complete.

See table 18 for an example of how the target in this example selects the rebuild and recalculate priorities using the rebuild/recalculate priority.

.

Table 18 - Rebuild priority selection

Codes	Description
00h	Rebuild or recalculate at least one stripe for every read or write request from an application client. (default)
01h	Suspend rebuild and recalculate operations during all read/write requests from any application clients.
02h-1Fh	Rebuild or recalculate at least 1/8th of a stripe for every read or write request from an application client.
20h-3Fh	Rebuild or recalculate at least 1/4th of a stripe for every read or write request from an application client.
40h-5Fh	Rebuild or recalculate at least 1/2th of a stripe for every read or write request from an application client.
60h-7Fh	Rebuild or recalculate at least one stripe for every read or write request from an application client.
80h-9Fh	Rebuild or recalculate at least two stripes for every read or write request from an application client.
A0h-BFh	Rebuild or recalculate at least four stripes for every read or write request from an application client.
C0h-DFh	Rebuild or recalculate at least eight stripes for every read or write request from an application client.
E0h-FEh	Rebuild or recalculate at least 16 stripes for every read or write request from an application client.
FFh	Do not accept any read/write requests from an application client until the rebuild or recalculate operation is complete.

Note:
If the redundancy group is configured as copy redundancy or S redundancy the target in this example will rebuild in 1 MByte stripes.

The contents of the NORMAL USER DATA TRANSFER SIZE field and the PERCENTAGE OF SEQUENTIAL READ TRANSFERS field in combination are used by the target to determine the amount of read ahead information to transfer from the disk drives.

See table 19 for an example of amount of read ahead the target in this example selects based on the contents of the NORMAL USER DATA TRANSFER SIZE field and PERCENTAGE OF SEQUENTIAL READ TRANSFERS field.

Table 19 - Read ahead selection

		PERCENTAGE OF SEQUENTIAL READ TRANSFERS			
		1% to 25%	25% to 50% or 0%	51% to 75%	greater than 75%
NORMAL USER DATA TRANSFER SIZE in logical blocks	1 to 128	1	2		3
	129 to 512 or 0	1	4		2
	greater than 512	1	5		6

Notes:
1) number of logical blocks to read ahead = no read ahead
2) number of logical blocks to read ahead = 2 x normal user data transfer size
3) number of logical blocks to read ahead = 4 x normal user data transfer size
4) number of logical blocks to read ahead = normal user data transfer size
5) number of logical blocks to read ahead = normal user data transfer size / 2
6) number of logical blocks to read ahead = normal user data transfer size up to max cache size

2 Example

A SCSI storage array is connected to an application client that runs applications that normally transfer 4 Kbytes of data at a time with sequential user data read requests occurring 60% of the time. In this example the application client issues a CREATE/MODIFY STORAGE ARRAY CONFIGURATION service action to the SCSI storage array with the NORMAL USER DATA TRANSFER SIZE field set to 8 blocks (each block is 512 bytes) and the PERCENTAGE OF SEQUENTIAL READ TRANSFERS field set to 60.

The SCSI storage array will create a volume set with a user data mapping as shown below:

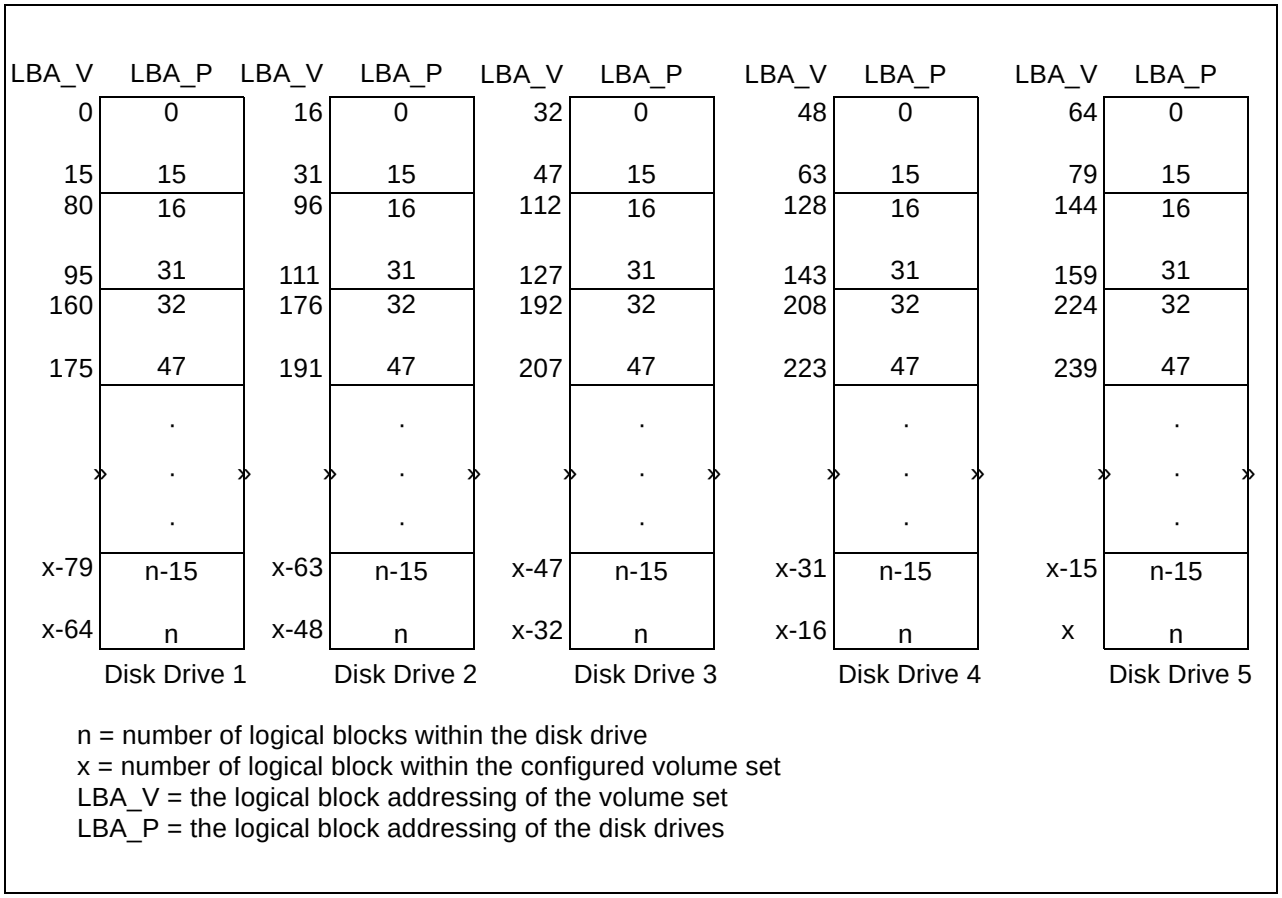

If, for example, the application client sends a read request for a 4 Kbyte transfer starting with LBA_V 112 then the storage subsystem would read LBA_Ps 16 through 23 from disk drive 3 and transfer that information to the application client. The SCSI storage array would also read ahead LBA_Ps 24 through 31 from disk drive 3 and LBA_Ps 16 through 23 from disk drive 4. The read ahead information is placed into the SCSI storage arrays' cache in anticipation that the next read from the application client will request that information.