

AssuredSAN 3xx4

Config Guide

Q3 2015



The following document provides basic specifications on the AssuredSAN 3xx4 (EX) series.

1.	WHAT'S NEW	4
2.	CONFIGURATION OVERVIEW:	5
3.	MODEL #'S	6
3.1.	Shipping Notes, Components, & Weights:	6
3.1.1.	Notes:	6
3.1.2.	Components/Cables:	6
3.1.3.	Weights:	7
4.	CONTROLLERS	9
4.1.	Protocols – CNC (FC&iSCSI) & SAS (Hybrid).....	9
4.2.	SFPs	9
4.3.	RAID Levels Supported	10
4.4.	Performance Overview	10
5.	CHASSIS:	11
5.1.	2U12 TLA Overview.....	11
5.2.	2U24 Overview	12
5.3.	How many RBOD/JBOD Chassis are Supported	13
5.4.	Hot Plug Expansion & Replacement Support.....	13
5.5.	Power Supplies/RAIL kits	13
6.	EXPANSION JBODS	14
6.1.	Expansion JBOD Overview	14
6.2.	Expansion JBOD Overview	Error! Bookmark not defined.
6.3.	JBOD Enclosure	14
6.4.	What JBODs are supported behind 3xx4 RBODs.....	15
7.	DRIVES	16
7.1.	What Drives Are Supported	16
7.2.	How many SSDs are supported (SFF in LFF Chassis)	16
7.3.	Drive Sleds & Air Management Sleds (AMS)	16
7.3.1.	Drive Sleds:	16
7.3.2.	Air Management Sleds:.....	16
7.4.	Encryption (SED/FDE)	Error! Bookmark not defined.
7.5.	512e Drives	17
8.	SOFTWARE / FEATURES.....	17
8.1.	Feature Scale Factors	17
8.2.	RealStor	19
8.2.1.	New User Interface:.....	19
8.3.	Storage Pools/Virtual Disk Groups/Volumes:.....	19
8.3.1.	Mapping in the V3 Interface:.....	20
8.3.2.	Provisioning and Mapping with the CLI Interface:.....	20
8.4.	Thin Provisioning:	21
8.5.	Performance Volumes LUNs and Rapid Rebuilds:	22
8.5.1.	Large LUNs:	22
8.5.2.	Rapid Rebuilds:	22
8.6.	Advanced Copy Services / Snapshots	23
8.6.1.	Differences between Copy on Write (COW) and Redirect on Write:	23
8.6.2.	Redirect on write snapshot overview:	23
8.7.	RealTier 2.0 and SSD read cache.	24
8.7.1.	Autonomic Sub-LUN Tiering:	Error! Bookmark not defined.
8.7.2.	Read Cache	24
8.7.3.	Key CLI Commands for RealTier 2.0 and Read Cache:	24
9.	WARRANTY.....	25
9.1.	General Warranty Options	25
9.1.1.	3 year warranty from date of shipment on all enclosures	25

9.2. Support Contact #'s / Resources	25
9.2.1. Customer Resource Center (CRC)	25
9.2.2. SANsolve Contact Information	25
9.2.3. Open a Support Ticket	26
9.2.4. Submitting RMA Requests	26
10. INTER-OP & AGENCY CERTIFICATIONS.....	27
10.1. Inter-Op Certifications:	27
10.2. Microsoft Certifications / Links:	Error! Bookmark not defined.
10.3. VMware Compatibility Guide:.....	Error! Bookmark not defined.
10.4. VMware Site Recovery Manager (SRM)	Error! Bookmark not defined.
10.5. Agency Certifications – Safety, Emissions, Harmonics, Country Approvals:	27
11. DATA SHEET SPECIFICATIONS – RBOD (POWER, TEMP, ACOUSTIC, DIMENSIONS).....	29
12. DATA SHEET SPECIFICATIONS – JBOD (POWER, TEMP, ACOUSTIC, DIMENSIONS)	30
13. GENERAL.....	31
13.1. NEBS level-3 certification	31
13.2. Super Capacitors	31
13.3. SimulCache™ Technology	31
13.4. Other Sources of Information	32

1. What's New

Dot Hill AssuredSAN 3004 Series Next-Generation, Entry-Level Storage Delivers Improved Price/Performance, Greater Capacity and High Speed Connectivity

3xx4 Controllers

The AssuredSAN 3004 models improve on previous-generation 3000 Series arrays by providing next-generation affordable storage to meet entry-level requirements with greater capacity and performance. In addition, 3000 Series customers can take advantage of easy field upgradability to the next-generation AssuredSAN 3004 RAID controller, simplifying deployment to achieve an instant performance and connectivity boost.

- *Performance*
 - Up to 3300 MB/s per second sequential reads, 2400 MB/s sequential writes
 - Performance for HD, 2K, 4K, and 3D workflows
 - Simultaneously stream to multiple hosts without disruption
- *Host Connections*
 - 2 ports per controller, 4 for dual controller models
 - 16Gb/8Gb Fibre Channel, 12Gb/6Gb SAS, and 10Gb/1Gb iSCSI host connectivity
- *Green Technology*
 - Patented EcoStor™ battery-free cache backup eliminating battery replacement.
- *Rugged, Mobile Applications*
 - AssuredSAN 3xx4 models are NEBS Level 3 and MIL STD 810G compliant to meet Telecommunications requirements and Military standards for shock and vibration.

New! GL210 Firmware

A new product firmware “GL210” is available for the AssuredSAN 3004 solutions. With GL210, the AssuredSAN 3004 solutions deliver up to 40,000 input/output operations per second (IOPS) from disk, sustained sequential read performance of up to 3300 MB/second and writes of up to 2400 MB/second. The Dot Hill RAID Companion Processor (RCP) allows 3004 solutions to deliver dramatic performance gains at very cost-effective price points.

New! RealStor Software for 3004:

RealStor™ is the next-generation storage operating system from Dot Hill which delivers application-aware intelligence for today's hybrid flash storage arrays. Easy-to-use RealStor features accelerate storage management operations to deliver data where customers need it, when they need it. With patent-pending innovations that deliver better efficiency, higher reliability, cost-effectiveness, and fast access for read data using SSD read-caching, RealStor is now available across Dot Hill's entire line of storage systems equipped with AssuredSAN 3004 storage controllers. The software features section of this configuration guide provides an overview of all RealStor features including the new user interface, thin provisioning, SSD read cache, and quick rebuilds.

New! 3004 SAS Controller:

Dot Hill AssuredSAN 3004 storage systems have a new addition to the family with the inclusion of 2 port 12Gb/6Gb SAS controller (4 ports per system). This SAS controller features an optional set of host-port expansion cables to allow up to 4 dual-port host connections. This controller is offered in addition to the existing 3004 controller types, which include 16Gb/8Gb Fibre Channel, and 10Gb/1Gb iSCSI with Dot Hill's unique, flexible Fibre Channel/iSCSI converged interface. With the ability to support up to 288 terabytes of total capacity, Dot Hill AssuredSAN 3004 storage systems are available in a variety of 2.5-inch and 3.5-inch hard disk drives (HDDs).

2. Configuration Overview:

As part of the AssuredSAN 3000 series, the AssuredSAN 3xx4 storage arrays are designed to deliver reliable, scalable performance for small to medium sized datacenter operations. They include energy saving drive-spin-down, EcoStor™ battery-free cache backup, data protection software to complement your backup and recovery plans. The AssuredSAN 3xx4 storage array provides:

Controllers:

- High performance controllers: 40,000 IOPs, 3300 MBps Read, 2400 MBps Write
- Active/Active controller architecture with two ports per controller
- 12GB system cache (6GB per controller)
- FC, iSCSI, or SAS interface

Chassis Options

The AssuredSAN 3xx4 is supported in two (2) industry leading chassis that can serve as an RBOD or JBOD.

Low Density	
2U12:	12 x 3.5" drives per chassis*
2U24:	24 x 2.5" drives per chassis

*Can also support 2.5" drives

JBOD Options

- Add the above chassis as a JBOD expansion for added performance and capacity. A maximum of three (3) JBOD expansion chassis behind a single RBOD is supported; (See JBOD section for details).

Drives Options

All AssuredSAN 3xx4 chassis support multiple drive types that are dual sourced from several manufacturers to provide quality and consistency of supply. The AssuredSAN 3xx4:

- Can be ordered with a single TLA that include one type of drives
- Can be ordered as a drive-less system so you can add multiple drive types
- Include industry leading capacity drives for SSDs, HDDs, and Self Encrypting Drives (SED).

SFP Options for long term flexibility

- CNC based controllers with modular SFPs including: 8Gb/16Gb FC and/or 1GbE/10GbE iSCSI connectivity
- 12Gb SAS connectivity (no SFP needed)

Base Software – RealStor included at no charge and includes:

- Enhanced User Interface – Includes both Traditional RAID (V2) and Virtualization mode (V3)
- RealCache™ – SSD Read-Cache
- RealThin™ Thin Provisioning
- RealQuick™ Rapid RAID Rebuild
- RealPool™ Virtualized pooling

Optional/Licensable Software - Advanced Copy Services / Data Management Software (DMS)

- Includes AssuredSnap, AssuredCopy, and Assured Remote software for no charge up to 90 days. After 90 days the software will become inactive without a license key which enables up to 1024 snapshots.
- RealSnap™ – When run in virtualized mode, the same software can provide up to 1024 virtualized redirect on write snapshots. Assured Copy is not needed, and Remote Replication features are expected to be available later this year.

Warranty:

- Hardware: Standard 3 year warranty, with advance exchange, from date of shipment on all enclosure FRUs.
- Software: 90 day free of defects.

3. Model #'s

RBOD (RAID) models are available with both converged network controller (CNC) and SAS Models.
JBOD (EX) models have 6G SAS ports to connect with the RBOD

CNC Connectivity	SAS Connectivity	JBODs
D3824 = 2U x 24 drives D3834 = 2U x 12 drives	D3524 = 2U x 24 drives D3534 = 2U x 12 drives	J6G24 = 24 bay J6G12 = 12 bay

3.1.Shipping Notes, Components, & Weights:

3.1.1. Notes:

SFPs: Each 4-pack will ship in a separate box
When ordering a part (PFRUKK14) you only order quantity of 1 – you will receive 4 SFPs.

Drives: When ordering bulk packs ending in 14 or 24, order the exact # of drives that are needed.
So ordering 22 x (PFRUKFXAXN161-24 / DD,SV3,900GB,10K,6G,SAS,BB,Bulk) = 22 x 900G drives.

For Low density chassis 2U12 and 2U24: Drives ship in the chassis

3.1.2. Components/Cables:

The following components are included with each system:

RBOD (12 bay)	
Qty	Descriptions
2	Power Cord IEC-C14 to C13 250V 10A Black 1.5m
1	Cable, USB A Male-to-Mini-B, 5-Wire, 1.5m, Black
1	Rack Mount Kit
1	Bezel (Bezel with filter is available)

JBOD (12 bay)	
Qty	Descriptions
2	Power Cord IEC-C14 to C13 250V 10A Black 1.5m
2	Cable miniSAS8088 to miniSAS8088 4x 0.6m,Gen
1	Rack Mount Kit
1	Bezel (Bezel with filter is available)

RBOD (24 bay)	
Qty	Descriptions
2	Power Cord IEC-C14 to C13 250V 10A Black 1.5m
1	Cable, USB A Male-to-Mini-B, 5-Wire, 1.5m, Black
2	Ears (Right and left side)
1	Rack Mount Kit
1	Bezel

JBOD (24 bay)	
Qty	Descriptions
2	Power Cord IEC-C14 to C13 250V 10A Black 1.5m
2	Cable miniSAS8088 to miniSAS8088 4x 0.6m,Gen
2	Ears (Right and left side)
1	Rack Mount Kit
1	Bezel

3.1.3. Weights:

System weights: Below are system weights for 3xx4 systems without packaging. These weights are approximate due to slight variations in drives (multiple vendors) and pounds listed have been rounded.

RBODs with dual-controller modules				
Item	2U24		2U12	
	Kg	Lb	Kg	Lb
Chassis	8.5	19	8.5	19
PSU x 2	3.8	8	3.8	8
RIOM x 2	4.4	10	4.4	10
SFF x 24	6	13	-	-
LFF x 12	-	-	9.3	20
Weight	22.7	50	25.9	57

JBODs with dual-expansion modules				
Item	2U24		2U12	
	Kg	Lb	Kg	Lb
Chassis	8.5	19	8.5	19
PSU x 2	3.8	8	3.8	8
JIOM x 2	2.9	6	2.9	6
SFF x 24	6	13	-	-
LFF x 12	-	-	9.3	20
Weight	21.2	47	24.5	54

Dot Hill has made ordering easy by creating Top Level Assembly (TLA) part #s. These TLAs ship with dual controllers, drive full, and AC power supplies, and are available in 2U12 and 2U24 chassis.

3xx4 - CNC (FC/iSCSI)		
	TLA	Short Description
CNC	D3824CA07215BA	3824,2RM,12GB,24x300SAS,15k,AC,V2,DH
	D3824CA14410BA	3824,2RM,12GB,24x600SAS,10k,AC,V2,DH
	D3824CA21610BA	3824,2RM,12GB,24x900SAS,10k,AC,V2,DH
	D3824CA28810BA	3824,2RM,12GB,24x1.2SAS,10k,AC,V2,DH
	D3824CN24007BA	3824,2RM,12GB,24x1TB SAS,7k,AC,V2,DH
	D3834CN24007BA	3834,2RM,12GB,12x2TB SAS,7k,AC,V2,DH
	D3834CN36007BA	3834,2RM,12GB,12x3TB SAS,7k,AC,V2,DH
	D3834CN48007BA	3834,2RM,12GB,12x4TB SAS,7k,AC,V2,DH
	D3834CN72007BA	3834,2RM,12GB,12x6TB SAS,7k,AC,V2,DH

3xx4 - (SAS)		
	TLA	Short Description
SAS	D3524CA07215BA	3524,2RM,12GB,24x300SAS,15k,AC,V2,DH
	D3524CA14410BA	3524,2RM,12GB,24x600SAS,10k,AC,V2,DH
	D3524CA21610BA	3524,2RM,12GB,24x900SAS,10k,AC,V2,DH
	D3524CA28810BA	3524,2RM,12GB,24x1.2SAS,10k,AC,V2,DH
	D3524CN24007BA	3524,2RM,12GB,24x1TB SAS,7k,AC,V2,DH
	D3534CN24007BA	3534,2RM,12GB,12x2TB SAS,7k,AC,V2,DH
	D3534CN36007BA	3534,2RM,12GB,12x3TB SAS,7k,AC,V2,DH
	D3534CN48007BA	3534,2RM,12GB,12x4TB SAS,7k,AC,V2,DH
	D3534CN72007BA	3534,2RM,12GB,12x6TB SAS,7k,AC,V2,DH

SFP's (1 x 4-Pack needed for 3004)		
Qty = 1 (4 packs) Select based on speed & type	PFRUKK11	SFP,1Gb Copper RJ45 -1D,4 Pack Set
	PFRUKK12	SFP,8Gb 850nm Rate Select,4 Pack Set
	PFRUKK13	SFP+,10GigE 3.3V 850nm LC Mod,4 Pack Set
	PFRUKK14	SFP+,16GFC 3.3V 850nm LC Mod, 4 Pack Set

Optional RealTier 2.0: Software & Support		
RealStor SW	SW-RealStor2	RealStor virtualization SW Download - includes support
Advanced Copy Services: Software & Support		
DMS SW	SW-DMS-R01-3K	Copy,Snap,Remote Pkg for AssuredSAN 3xx4
DMS Support	SW-DMS-R01-3K-M1	DMS,1-Yr Maint Ext,AssuredSAN 3xx4
Hardware & Support		
Hardware Support	DS-N-R-1Y-A-SS	NBD Adv Ex from Depot, RAID, 1 YEAR, AME

Port Expansion Cables		
Single and 4 packs available	PFRUKL23-01	Cbl Pkg,Y HDminiSAS/HDminiSAS,2m,MLX
	PFRUKL23-04	Cbl Pkg,Y HDminiSAS/HDminiSAS,2m,MLX,4BK
	PFRUKL24-01	Cable Pkg,Y HDminiSAS/miniSAS,2m,MLX
	PFRUKL24-04	Cable Pkg,Y HDminiSAS/miniSAS,2m,MLX,4BK

The 3xx4 can be ordered a la carte and additional TLAs (chassis, drives, components, software) are listed in respective sections. For a complete list, check with your local sales person, value-added reseller, or distribution partner.

4. Controllers


4.1. Protocols – CNC (FC&iSCSI) & SAS

AssuredSAN 3xx4 controller architecture maximizes performance, including
Two host ports per controller

- 3xx4 controller supports 8Gb FC, 16Gb FC, 1GbE iSCSI or 10GbE iSCSI SFPs.
- 3xx4 SAS controller supports 6Gb and 12Gb SAS host connectivity using mini-SAS HD Cables. (Also supports four port-expansion (Y-cables) to enable up to four dual-port hosts to be attached.)

6GB transportable read/write cache per controller.

Battery-free cache backup with super capacitors and compact flash



- | | |
|--|--|
| 1. CNC ports (0,1) shown with FC SFPs | 5. Network Port |
| 2. CLI port | 6. Service Port 1 (used by service personnel only) |
| 3. Service port 2 (used by service personnel only) | 7. Disable button (used by engineering only) |
| 4. Reserved for future use | 8. Mini-SAS expansion port |

- **Single Controllers:**
Single controllers are supported in the 2U12 and 2U24 but NOT recommended due to lack of failover capability.

4.2. SFPs

For CNC versions (D38xxx), one four-pack set is advised. Both FC and iSCSI transceivers contain unique identification preventing unqualified SFP's from being used in the 3xx4 (EX) controller.

Dot Hill branded SFP's must be ordered with the AssuredSAN 3xx4 controllers with the CNC (Converged Network Controller). They are available in 4 packs, so quantity 1 is recommended. They are available in 1G iSCSI, 10G iSCSI, 8G FC, or 16G FC. SAS and JBOD systems do not require SFP's.



TLAs for the SFPs for the 2U12 and 2U24 are below.

- | | |
|------------|--|
| • PFRUKK11 | SFP,1Gb Copper RJ45 -1D,4 Pack Set |
| • PFRUKK12 | SFP,8Gb 850nm Rate Select,4 Pack Set |
| • PFRUKK13 | SFP+,10GigE 3.3V 850nm LC Mod,4 Pack Set |
| • PFRUKK14 | SFP+,16GFC 3.3V 850nm LC Mod, 4 Pack Set |

4.3.SAS Port Expansion Cables

The SAS controller has 2 physical SAS ports which can accommodate an optional host port-expansion cable. Up to 4 of these cables can be connected to the SAS ports enabling 4 dual-port host connections. Each host-port expansion splits the 4 SAS lanes per port, so that 2 lanes are connected to each host port.

TLAs for the Host Port Expansion cables for the 2U12 and 2U24 are below:

- PFRUKL23-01 Cbl Pkg,Y HDminiSAS/HDminiSAS,2m,MLX
- PFRUKL23-04 Cbl Pkg,Y HDminiSAS/HDminiSAS,2m,MLX,4BK
- PFRUKL24-01 Cable Pkg,Y HDminiSAS/miniSAS,2m,MLX
- PFRUKL24-04 Cable Pkg,Y HDminiSAS/miniSAS,2m,MLX,4BK

4.4.RAID Levels Supported

The 3xx4 controller supports RAID 0, 1, 3, 5, 6, 10, 50

In addition to the usual RAID levels, the 3xx4 features several important additional levels. RAID 6 offers the highest level of RAID protection. It allocates two sets of parity data across drives and allows simultaneous write operations. It can withstand two simultaneous drive failures without downtime or data loss. RAID 10 is mirroring and striping without parity and allows large Disk Groups to be created with high performance and mirroring for fault tolerance. RAID 50 combines the block striping and parity of RAID 5 with the straight block striping of RAID 0, yielding higher performance than RAID 5 through the addition of RAID 0, particularly during writes.

4.5.Performance Overview

Model #	3xx4 (EX)
Firmware	GL210
# of ports - System	4
IOPS	40 K
Max Read - MBps	3300
Max Write - MBps	2400
Interface (per controller)	2 x 8G or 16G FC or 2 x 1G or 10G iSCSI or 2 x 12G or 6G SAS
Cache – per system	12 GB
GUI	RAIDar (v2) Storage Management Console (v3)
Value Add Software	DMS / RealStor
Max # of initiators	1024
Max # of LUNS	1024
Maximum # of chassis:	Up to 4
Max # of Drives/System:	2U12 = 48 2U24 = 96

5. Chassis:

5.1.2U12 TLA Overview



To build a system a la carte, choose from the TLAs below.

Note: Additional TLAs are available. For a full up-to-date list, see our price book.

Drive-less RBOD System (include rail kit):

D3834C000000BA	3834,2RM,12GB,No drives,AC,V2,DH	CNC
D3534C000000BA	3534,2RM,12GB,No drives,AC,V2,DH	SAS
D3834C000000BD	3834,2RM,12GB,No drives,DC,V2,DH	DC Power

Drive-less JBOD System (include rail kit):

J6G12C000000BA	6GJBOD,2JM,2U12,No drives,AC,V2,DH
J6G24C000000BA	6GJBOD,2JM,No drives,AC,V2,DH

Drives (sold as singles)

PFRUKFXGXN084-01	DD,LV2,2.0TB,7k,6G,N-SAS,BB,PKG
PFRUKFXGXN083-01	DD,LV2,3.0TB,7k,6G,N-SAS,BB,PKG
PFRUKTXGXN119-01	DD,LV2,4.0TB,7k,6G,N-SAS,BB,PKG
PFRUKTXGXN132-01	DD,LV2,4.0TB,7K,6G,ENCR,N-SAS,BB,PKG
PFRUKTXGXN194-01	DD,LV2,6.0TB,7K,6G,N-SAS,BB,PKG
PFRUKFXGXN150-01	AMS,LV2 Drive Blank, BB,PKG

SFF Drives (Single Pack) using a 3.5" Hybrid Form Factor drive sled

PFRUKRXGXN227-01	SSD,HV2,1.6TB,12G,SAS,BB,PKG
PFRUKRXGXN228-01	SSD,HV2,800GB,12G,SAS,BB,PKG
PFRUKRXGXN229-01	SSD,HV2,400GB,12G,SAS,BB,PKG
PFRUKTXGXN231-01	DD,HV2,1.2TB,10K,6G,ENCR,SAS,BB,PKG
PFRUKTXGXN230-01	DD,HV2,1.2TB,10K,6G,SAS,BB,PKG

Drives (sold as 24 packs)

PFRUKFXGXN084-24	DD,LV2,2.0TB,7k,6G,N-SAS,BB,24 BULK PK
PFRUKFXGXN083-24	DD,LV2,3.0TB,7k,6G,N-SAS,BB,24 BULK PK
PFRUKTXGXN119-24	DD,LV2,4.0TB,7k,6G,N-SAS,BB,24 BULK PK
PFRUKTXGXN132-24	DD,LV2,4.0TB,7K,6G,ENCR,N-SAS,BB,24 BULK
PFRUKTXGXN194-24	DD,LV2,6.0TB,7K,6G,N-SAS,BB,24BULK
PFRUKFXGXN150-48	AMS,LV2 Drive Blank,BB,48 Bulk Pack

SFPs

PFRUKK11	SFP,1Gb Copper RJ45 -1D 4pk
PFRUKK12	SFP,8Gb 850nm Rate Select,4pk
PFRUKK13	SFP+,10GigE 3.3V 850nm LC Mod,4pk
PFRUKK14	SFP+,16GFC 3.3V 850nm LC Mod, 4pk

Host Port-Expansion Cable

PFRUKL23-01	Cbl Pkg, Y HDminiSAS/HDminiSAS,2m,MLX
PFRUKL23-04	Cbl Pkg, Y HDminiSAS/HDminiSAS,2m,MLX,4BK
PFRUKL24-01	Cable Pkg, Y HDminiSAS/miniSAS,2m,MLX

PFRUKL24-04	Cable Pkg, Y HDminiSAS/miniSAS,2m,MLX,4BK
-------------	---

Controller/Miscellaneous components

PFRUKC61-01	38X4,CNC,2P,1RM-EX,FRU,PKG
PFRUKC76-01	35X4,SAS,2P,1RM-EX,FRU,BB,PKG
FHDW018-02	Rack Mount Kit,Shelf,Long,All HW 25"-36"
PFRUKA17-01	Chas+MP,2U12,6G,V2,BB,FRU,PKG
PFRUKA19-01	Chas+MP,2U24-S2,6G,V2,BB,FRU,PKG

5.2.2U24 Overview

To build a system a la carte, choose from the TLAs below.

Note: Additional TLAs are available. For a full up-to-date list, see our price book.

Drive-less RBOD System (include rail kit):

D3824C000000BA	3824,2RM,12GB,No drives,AC,V2,DH	CNC
D3524C000000BA	3524,2RM,12GB,No drives,AC,V2,DH	SAS
D3824C000000BD	3824,2RM,12GB,No drives,DC,V2,DH	DC Power

Drive-less JBOD System (include rail kit):

J6G12C000000BA	6GJBOD,2JM,2U12,No drives,AC,V2,DH
J6G24C000000BA	6GJBOD,2JM,No drives,AC,V2,DH

Drives (sold as singles)

PFRUKT76-01	DD,SFF,1.0TB,7K,6G,SAS,BB,PKG
PFRUKTXSEN113-01	DD SFF 600GB 10K 6G ENCR SAS BB, PKG
PFRUKF64-01	DD,SFF,600GB,10K,6G,SAS,BB,PKG
PFRUKT79-01	DD,SFF,900GB,10K,6G,SAS,BB,PKG
PFRUKTXSXN129-01	DD,SFF,1.2TB,10K,6G,SAS,BB,PKG
PFRUKTXSXE170-01	DD,SFF,1.2TB,10K,6G,SAS,ENCR,BB,PKG
PFRUKT81-01	DD,SFF,300GB,15K,6G,SAS,BB,PKG
PFRUKTXSXN207-01	DD,SFF,600GB,15K,6G,SAS,BB,PKG
PFRUKRXSXN108-01	SSD,SFF,400GB,SAS,BB,PKG
PFRUKRXSXN125-01	SSD,SFF,800GB,SAS,BB,PKG
PFRUKRXSXN215-01	SSD,SFF,1.6TB,6G,SAS,BB,PKG
PFRUKF31-01	AMS,SFF Drive Blank,BB,FRU,PKG

Drives (sold as 24 packs)

PFRUKT76-24	DD,SFF,1.0TB,7K,6G,SAS,BB,24 BULK PK
PFRUKTXSEN113-24	DD,SFF,600GB,10K,6G,SAS,ENCR,BB,24 BULK PK
PFRUKF64-24	DD,SFF,600GB,10K,6G,SAS,BB,24 BULK PK
PFRUKT79-24	DD,SFF,900GB,10K,6G,SAS,BB,24 BULK PK
PFRUKTXSXN129-24	DD,SFF,1.2TB,10K,6G,SAS,BB,24 BULK PACK
PFRUKTXSXE170-24	DD,SFF,1.2TB,10K,6G,SAS,ENCR,BB,24 PACK
PFRUKT81-24	DD,SFF,300GB,15K,6G,SAS,BB,24 BULK PK
PFRUKTXSXN207-24	DD,SFF,600GB,15K,6G,SAS,BB,24 BULK PK
PFRUKF31-48	AMS,SFF Drive Blank,BB,48 Bulk Pack

SFPs

PFRUKK11	SFP,1Gb Copper RJ45 -1D 4pk
PFRUKK12	SFP,8Gb 850nm Rate Select,4pk

PFRUKK13	SFP+,10GigE 3.3V 850nm LC Mod,4pk
PFRUKK14	SFP+,16GFC 3.3V 850nm LC Mod, 4pk

Host Port-Expansion Cable

PFRUKL23-01	Cbl Pkg,Y HDminiSAS/HDminiSAS,2m,MLX
PFRUKL23-04	Cbl Pkg,Y HDminiSAS/HDminiSAS,2m,MLX,4BK
PFRUKL24-01	Cable Pkg,Y HDminiSAS/miniSAS,2m,MLX
PFRUKL24-04	Cable Pkg,Y HDminiSAS/miniSAS,2m,MLX,4BK

Controller/Miscellaneous components

PFRUKC61-01	38X4,CNC,2P,1RM-EX,FRU,PKG
PFRUKC76-01	35X4,SAS,2P,1RM-EX,FRU,BB,PKG
FHDW018-02	Rack Mount Kit,Shelf,Long,All HW 25"-36"
PFRUKA17-01	Chas+MP,2U12,6G,V2,BB,FRU,PKG
PFRUKA19-01	Chas+MP,2U24-S2,6G,V2,BB,FRU,PKG

5.3.How many RBOD/JBOD Chassis are Supported

Up to 1 RBOD + 3 JBOD are supported.

5.4.Hot Plug Expansion & Replacement Support

All 3xx4 models support hot plug expansion and replacement of redundant controllers, enclosures, fans, power supplies, and I/O modules for simple, fast installation and maintenance. Hot add expansion of disk enclosures is also supported.

5.5.Power Supplies/Rail KitsPower Rating (RBODs)

2U12/2U24	595W supply	87% for 50% load - designed to meet ENERGY STAR compliance
-----------	-------------	--

Power Switch: DHS has released a second source power supply for use in 2u12 and 2u24 products. This new power supply from Acbel is available with or without a power switch and is a direct replacement option for the current Emerson model supplies. See Tech bulletin: 81-00000086-00-01

AC Power Supply

2U12/2U24:	PFRUKE08-01	Power Supply,595W,AC,2U12/24,FRU,PKG
------------	-------------	--------------------------------------

DC Power Supply

2U12/2U24:	PFRUKE02-01	Power Supply,575W,DC,2U12/24,FRU,PKG
------------	-------------	--------------------------------------

Rail Kits:

2U12/24	FHDW018-02	Rack Mount Kit,Shelf,Long,All HW 25"-36"
---------	------------	--

6. Expansion JBODs

6.1.Expansion JBOD Overview

Just as the user has a choice of chassis for the array head (LFF and SFF drive bays, AC or DC powered), so also do they have a choice of expansion disk enclosures accommodating either drive size. The J6G12 and J6G24 can be hot-added to an operating array. SFF and LFF Array RBODs and JBODs are mixed with few limitations.

J6G12, 3.5-inch Disk Enclosure. This 2U unit has twelve LFF (3.5-inch) drive bays and accepts dual-ported SAS, SEDs and SAS NL drives. The pre-configured LFF Drive Enclosures have two I/O modules and support both single and dual controller arrays.

- This 3.5-inch disk enclosure can be attached to any 3xx4 RBOD.
- Each configured model ships standard with two .5m mini-SAS to mini-SAS cables for connection to the AssuredSAN array expansion port or existing disk enclosure cascade port.
- LFF and/or SFF Disk Enclosures can be mixed up to the maximum of 4 total Disk Enclosures
- The J6G12 does not support LFF SATA HDDs.


J6G24, 2.5-inch Disk Enclosure. This 2U storage enclosure is designed to support twenty four 2.5-inch Universal form factor (SFF) 12Gb, SSD, SAS, SEDs or SAS NL hard drives. It ships standard with dual I/O modules installed.

- This 2.5-inch disk enclosure can be attached to any 3xx4 LFF or SFF RBODs.
- Each configuration ships with a two .5m mini-SAS to mini-SAS cables for connection to the AssuredSAN 3xx4 array expansion port or existing disk enclosure cascade port.
- LFF and/or SFF Disk Enclosures can be mixed up to the maximum of 4 total disk Enclosures.
- The J6G24 does not support SFF SATA HDDs.

Scalability: The AssuredSAN 3xx4 array configurations are designed to allow an installation to begin with smaller capacity and be able to grow gradually as needed. The flexibility of SSD, SAS or SAS NL drive technology, form factors, sizes, speeds, and costs per GB allows a system to easily fit in almost any budget.

6.2.JBOD Enclosure

AssuredSAN J6Gxx standard disk enclosure (2U) - rear panel components:



- | | |
|--------------------|-----------------|
| 1. Power Indicator | 4. SAS In Port |
| 2. Fault Indicator | 5. SAS Out Port |
| 3. Unit Locator | |

6.3.What JBODs are supported behind 3xx4 RBODs

	RBOD	Line	AssuredSAN	
		Model	3004	
		Bays	2U12	2U24
		TLA	38x4/35x4	
			If you have the above RAID product, the below are "recommended" JBOD's	
Titanium IOM	24	3120	AssuredSAN 3000 "Titanium" JBODs behind LX RBODs are supported, but LX JBODs are recommended.	
	12	3130		
	24	4120		
	12	4130		
Gallium IOM	24	4124	J6G24	J6G24
	12	4134	J6G12	J6G12
	48	4144	no	no
	56	4154	no	no

Driveless JBOD (AC power) TLAs:

2U12: J6G12C000000BA 6GJBOD,2JM,2U12,No drives,AC,V2,DH
 2U24: J6G24C000000BA 6GJBOD,2JM,2U24,No drives,AC,V2,DH

7. Drives

7.1.What Drives Are Supported

The following table shows what drives are supported in what chassis.

Supported Drives (Q2 2015)			AssuredSAN 3004	
			2U12	2U24
LFF	7k	1TB	Yes*	
		2TB	Yes	
		3TB	Yes	
		4TB	Yes	
		6TB	Yes	
		8TB	Yes	
SFF	10K	300GB	EOL	
		600GB		Yes
		900GB		Yes
		1.2TB	Yes*	Yes
		1.8TB		Yes
	15k	300GB		Yes*
		450GB		Yes*
		600GB		Yes*
	7k	1TB		Yes
		2TB		Yes
	SSD	200GB		Yes
		400GB		Yes
		800GB	Yes*	Yes
		1.6TB	Yes*	Yes

Key Yes* = but not on price book

7.2.How many SSDs are supported (SFF in LFF Chassis)

2U12: Up to 12 SSDs or 2.5" HDDs are supported

2U24: Up to 24 SSDs

7.3.Drive Sleds & Air Management Sleds (AMS)

7.3.1. Drive Sleds:

Below are drive sleds that house drives for AssuredSAN 3004 chassis.

2U12:	PFRUKF89-12	Sled,LV2,Kit w/Screws,BB,12 Bulk Pack
2U24:	PFRUKF88-24	Sled,SFF,Kit w/Screws,BB,24 Bulk Pack

7.3.2. Air Management Sleds:

2U12:	Single:	PFRUKFXGXN150-01	AMS,LV2 Drive Blank, BB,PKG
	48 Pack	PFRUKFXGXN150-48	AMS,LV2 Drive Blank,BB,48 Bulk Pack

2U24:

Single:
48 Pack:PFRUKF31-01
PFRUKF31-48AMS,SFF Drive Blank,BB,FRU,PKG
AMS,SFF Drive Blank,BB,48 Bulk Pack

7.4.512e Drives

For AssuredSAN storage arrays, Dot Hill supports emulation, or 512e drives, rather than the 4K native sector format for the drives. These drives include emulation (e) which supports backward compatibility for the many applications and legacy operating systems that do not support 4K native drives. Dot Hill continues to work actively with its drive suppliers and firmware teams to optimize drives to reduce the potential performance impacts that come with emulation and enable your organization to gain greater value from 512e drives.

Older versions of application software such as virtualization software that reside between the operating system and your storage firmware may not fully support 512e drive emulation. If not, performance degradation might result. Please ensure that you have upgraded to the most recent version of any software that might be affected, and refer to its documentation for further information.

8. Software / Features

8.1.Feature Scale Factors

	3004 (EX)	
	Traditional RAID	Virtual
RAID stripe chunk sizes (min/max)	16/512	Auto (512KB-64KB)
Drives per RAID: 0 (min/max)	2/16	2/2 (Read cache only)
Drives per RAID: 1 (min/max)	2/2	2/2
Drives per RAID: 3 (min/max)	3/16	N/A (Not allowed for virtual disk groups)
Drives per RAID: 5 (min/max)	3/16	3/16
Drives per RAID: 6, 10 (min/max)	4/16	4/16
Drives per RAID: 50 (min/max)	6/32	N/A (Not allowed for virtual disk groups)
Max dedicated spares per disk group	4	N/A
Max global spares per system	16	16
Maximum TBs per LUN	128	128
Max TBs per Virtual Pool	N/A	329.5TB per pool w/ snaps, 2PB per pool w/ snaps disabled
Max Chassis per system	4	4
Drives per system, 2U24 array plus max 2U24	96	96
Drives per system, 2U12 array plus max 2U12 JBODs	48	48
Disk Groups per system	64	32
Pools per system	64	2 (1 per controller)
Initiators per port	1024	1024
Initiators per system	8192	8192
Commands (preferred path) per LUN, 4 port controller (System limit)	4096	4096
Max Queue Depth per port (SAS)	1024	1024
Max Queue Depth per port (iSCSI)	1024	1024
Max Queue Depth per port (FC)	1024	1024

Mappable Physical Volumes per Pool	128	N/A
Maximum LUNs per system (This means mapped volumes)	1024	1024
Maximum Volumes per system	512 *note1	2048
Snapshots per system (net usable)	512 *note2	1024 (EX only with license)
Master Volumes per system	128 *note4	1024 base volumes
Snapshots per Master Volume	127	254 snaps per base (EX only with license)
Number of replication sets per Linear Disk Group	4 *note5	N/A
Number of replication snapshot per Replication Set	127	N/A
Max replication snapshots per system	512	N/A
Remote snap limits: Sets (S)	16	N/A
Remote snap limits: Nodes (N)	2	N/A
Remote snap limits: Remote Systems (RS)	3	N/A
Remote snap limits: Chains (C)	1	N/A
Max snapshot-related volumes (snaps + RSR snaps)	1024	1024 (note: RSR applies to linear only)
Max snaps per Pool	127 (DMS snaps)	512 (PS snaps)
Max Disk Group size	128 TB	256 TB design limit
Max SCSI reservations (system)	1024	1024
Max SCSI reservations (LUN)	1	1
Max SCSI registrations (system)	32768	32768
Max SCSI registrations (LUN FC or SAS)	128	4096
Max SCSI registrations (LUN iSCSI) – dependent on IQN length	85 to 128	4096
Max Initiators per Host	128	128
Max Hosts per Host Group	256	256
Max Host Groups per System	32	32
Max Volumes per Volume Management Group	1024	1024
Max Volume Management Groups per System	256	256

Note:

1. For "Maximum Volumes per system 512" under linear mode, it is really up to 1024, but 512 recommended.
2. For "Snapshots per system (net usable) 512" under linear mode, it is really up to 1016, but 512 recommended.
3. GL200 (and GL1xx) allows up to 1024 shared between volumes and snapshots.
4. For "Master Volumes per system" under linear mode, it is really up to 1024, but 128 recommended.
5. For "Number of replication sets per Vdisk" under linear mode, it is really up to 16, but 4 recommended.

8.2.RealStor

RealStor gives users access to several new key features including thin provisioning with UNMAP, SSD read cache, storage pooling, fast rebuilds and virtual pool provisioning. The new features will be accessed via a new User Interface, referred to as Storage Management Console V3. This new management GUI brings a new modern look and feel to array management. Existing customers can choose to use the new V3 or to continue to use the previous generation V2 if new virtualization features are not required.

8.2.1. New User Interface:

RealStor introduces advanced storage software for hybrid arrays, built for easy to use storage management, and optimizing performance and capacity needs to address today's dynamically changing workloads. The new .v3 Storage Management Console for RealStor provides all the capabilities of the prior interface with an updated look and feel that scales well for managing large numbers of volumes and hosts. New systems default to the new .v3 interface.

v3: Specifies to use the new WBI interface to manage virtual and linear storage, and to use terminology in CLI output and system messages that is oriented to managing a hybrid (linear and virtual) system, such as “disk group” for disk groups and “pool” for pools. This is the default for a new installation.

8.3.Storage Pools/Virtual Disk Groups/Volumes:

Storage Pools allow data on a given LUN to span across all drives in a pool. When capacity is added to a system, the user is also getting a performance benefit of the additional spindles –hence the term Wide Striping. Storage Pools are available with GL200 firmware or newer.

Virtual Disk Groups

When creating a Virtual Disk Group, the user can specify the RAID level (RAID 1, 5, 6, 10) and which Pool the Disk Group belongs to. Disk Groups must contain like disk types: SSD, Enterprise SAS, NL SAS. The user must select a valid group of available HDDs or SSDs for the chosen RAID type. These configuration rules are the same as creating a VDisk with the prior V2 interface. Unlike creating VDisks in the V2 interface, the chunk size is automatically set to 512KB and does not need to be modified by the user. An example of the V3 Add Disk Group Screen is shown in Figure 1 below:



Figure 1 – Add Disk Group Screen

Volumes in Virtual Disk Pools:

To further explain the virtual disk pool concept, we will first look at how volumes are treated in the virtual disk pool. One of the key concepts of the RealStor firmware is the addition of a virtualization layer which allows a volume's data to be segmented into discrete 4MB entities called pages. These 4MB pages are managed individually, and can allow a volume's data to be spread over potentially multiple underlying storage pool Virtual Disk Groups and potentially multiple tiers (SSD, Enterprise SAS, NL SAS) of disks.

8.3.1. Mapping in the V3 Interface:

Once the users have added their Disk Groups to create the Pools, they may now create and map volumes. This is very similar to how the customers currently create volumes in the V2 WBI, however volumes are now assigned to a pool instead of a VDisk.

RealStor includes the ability to assign volumes to groups. This grouping capability enables large numbers of initiators to be mapped to LUNs in a single operation – improving the efficiency and organization of host configuration. In addition, the V3 interface incorporates search filters that greatly simplify the task of managing large numbers of volumes and reduce management time.

To map a volume, the user may select the Mapping topic, or click on the volume and select Mapping under the action tab. The Mapping screen in the V3 interface is different than the V2 version and contains several enhancements.

Users may now create initiator names, assign to Hosts, and then assign Host to Host Groups. This greatly simplifies the tasks of managing and mapping large numbers of Hosts that may have the same mapping requirements. Filter capabilities also exist for the Host and Host group tables.

RealStor allows for users to do default mapping by selecting “All other initiators.” Selecting this will make the volume available to all host ports. The user may also do explicit mapping - mapping to individual initiators, hosts, or host groups. The Storage Management Console will map to both the A and B controller by default for each port. Figure 2 is an example of the Mapping Screen.




Figure 2 – Mapping Screen in RealStor’s Management Console

8.3.2. Provisioning and Mapping with the CLI:

RealStor will support a robust CLI. RealStor maintains a set of commands for the Linear (V2) interface, and a separate set of commands specific to the new Virtual (V3) interface. While much of the command-set is shared between the two, there are some significant differences which are required to maintain backwards compatibility while allowing users to manage both versions in the same CLI.

While this document is not intended to highlight all the differences, some of the commands specific to provisioning and mapping that customers will want to review are highlighted here. For detailed description of all the changes, CLI Reference guide will be available.

First, specific to setting a management interface, **set protocols** and **set cli-parameters** will include support for setting the management-mode interface to **v2** or **v3** at both a global (set protocols) or session level (set cli-parameters). This setting will dictate which CLI commands are used.

With the introduction of the term **disk group**, RealStor will provide a new set of CLI commands to configure disk groups: **add disk-group**, **remove disk-groups**, **show disk-groups**, etc. **Disk groups** are a group of disks that together form a RAID set.

In a **V2 (RAIDar 2.0)** environment, disk groups are exactly the same as vdisks. Commands used to manage vdisks in GL105 code will be the same as in the **V2** interface.

With volumes now residing in **pools**, a new set of CLI commands is available to manage them; **show pool**, **remove pool**, etc.

As highlighted above, the CLI will allow supporting for **host groups** in both the **V2** and **V3** interface modes. However, the term host can have a different meaning depending on whether you are grouping, or using the legacy context of the command. RealStor will also introduce **volume groups** for virtual volumes. For additional information regarding CLI command changes, please refer to the CLI Reference Guide for a more detailed description.

8.4. RealThin: Thin Provisioning

Thin Provisioning allows storage allocation of physical storage resources only once they are consumed by an application. Thin Provisioning also allows over-provisioning of physical storage pool resources allowing ease of growth for volumes without predicting storage capacity upfront. Thin Provisioning is available with GL200 firmware or newer.

Thin Provisioning and Modern File Systems

Many modern operating systems become automatically aware of an array's ability to support thin provisioning. This is true of Windows 2012 / 2012 R2 and Linux with newer 6.x kernels. When using a non-thin provisioned storage array the deletion of a file does nothing more than rearrange the file system pointers.

If the array supports thin provisioning, as with RealStor, thin-aware operating systems will notify the storage array to free the storage that had been allocated to the file that was deleted. This happens with the SCSI UNMAP command. This will allow for the re-use of the capacity that had been allocated for the deleted file.

Windows 2012 / 2012 R2 also issues notifications when the allocations reach a high percentage of the "advertised" capacity. This, along with the capacity usage notifications, signals the user that it may be time to add more physical storage.

How Do I Enable Storage Over-Commitment With RealStor?

Overcommit is a pool setting (default on). When enabled, the aggregate capacity of a volume can exceed the physical backing storage.

The user can create a volume on a pool with over-commitment Enabled. No configuration needs to be done to create such a volume. Figure 3 shows that a volume can be logically larger than the pool that it's assigned to.




Figure 3 – Volume that is logically larger than the pool

Note: If you try to disable overcommit and the total space allocated to thin-provisioned volumes exceeds the physical capacity of their storage pool, an error will say that there is insufficient free disk space to complete the operation and overcommit will remain enabled.

Thresholds and Notifications

Users with a manage role can view and change settings that affect the thresholds and corresponding notifications for each storage pool.

- Low Threshold – When this percentage of pool capacity has been used, Informational event 462 is generated to notify the administrator. This value must be less than the Mid Threshold value. The default is 25%.
- Mid Threshold – When this percentage of pool capacity has been used, event 462 is generated to notify the administrator to add capacity to the pool. This value must be between the Low Threshold and High Threshold values. The default is 50%. If the commitment setting is enabled, the event has Informational severity; if the over-commitment setting is disabled, the event has Warning severity.
- High Threshold – When this percentage of pool capacity has been used, Critical event 462 is generated to alert the administrator that it is critical to add capacity to the pool. This value is automatically calculated based on the available capacity of the pool minus reserved space.

8.5. Performance Volumes/LUNs and Rapid Rebuilds:

8.5.1. Large LUNs:

The RealStor firmware release introduces the concept of Virtual Disk Pools – which are comprised of multiple Virtual Disk Groups (a disk group is similar to a VDisk). This now decouples the LUN from being restricted to a single VDisk. A Volume's data on a given LUN can now expand across all disk drives in a pool. When capacity is added to a system, the user is also getting a performance benefit of the additional spindles –hence the term Performance Volume.

- General “Best Practice” for VMFS datastores is to use large LUNs for use by multiple VMs. Large LUNs gives VMware vSphere users the most flexibility by not requiring storage administrators to provision new storage every time a new VM is created. Large LUNs also minimize rescans for new LUNs on the ESX server, which can impact VMFS I/O.
- Operating Systems now support up to 16TB file systems.
- As enterprise workloads for virtual environments grow in size and in performance demands, virtual hard disk (VHD) formats need to accommodate them. Hyper-V in Windows Server 2012 introduces a new version of the VHD format called VHDX which supports virtual disks up to 64TB in size.

RealStor supports large, flexible Volumes with sizes up to 128TB spanning up to 96 HDDs (architectural limit) providing seamless capacity expansion. Data automatically reflows to new drives and IOPs scale with additional storage.

8.5.2. RealQuick: Rapid Rebuilds

Rebuild of conventional (Linear) storage requires that the entire VDisk is rebuilt prior to returning the VDisk (and thus the Volumes on the VDisk) to a fault tolerant state. The rebuild operates at the VDisk level and has no knowledge of Volumes or customer data contained therein.

A customer could have a single small Volume on a large VDisk. During the rebuild, even after it has completed the stripes where the Volume resides, the VDisk remains in a degraded or critical state. During this period, another drive failure would take the VDisk from degraded to critical, or from critical to down. In the latter case, the customer may have lost data. In the prior case, the customer's data becomes at risk once again, even though that area of the VDisk had been rebuilt.

In another example, the customer may have created a Volume or Volumes consuming the entire VDisk space, but may have actually allocated very little. Since linear storage has no means to know what stripes contain customer data and

which do not, that small amount of data will remain at risk until the entire VDisk has been rebuilt. Also, performance suffers during the entire rebuild.

Another exposure is encountering a media error. If the array is rebuilding two drives and encounters a media error on a third, the rebuild will fail. This leaves the VDisk in a critical state, exposed, and with no way to rebuild the VDisk. It is especially unfortunate if this media error resides in unallocated space. The outcome is the same, linear storage must assume all stripes contain customer data.

With the introduction of Virtual Storage, we have a new set of technologies to improve these conditions. RealStor enables Virtual Disk Pools which track how much data has been allocated for each Volume and where it resides in the Pool, at a granularity of 4 MB. The Virtual Pool algorithm decides where to allocate each new 4 MB page.

Advantages of Quick Rebuild:

- LUNs are spread across many RAID sets, so one RAID set rebuilding only affects a fraction of all disk I/Os.
- Less rebuild work to do – so affected disks can return to full performance quickly.
- Volumes become fault tolerant more quickly – user data is restored to full protection more quickly.
- The rate of a RAID rebuild is directly proportional to the amount of unused disk space in a Virtual Disk Pool.

8.6. RealSnap: Advanced Copy Services / Snapshots

Unlike prior AssuredSAN arrays which use VDisks, RealStor users will create storage pools by creating and adding Virtual Disk Groups into a Virtual Disk Pool. The new copy services will run within these pools and leverage the new storage page design. RealStor will allow the existing COW snapshots and the new Virtual Snapshots for Virtual Disk Pools to coexist on the same system.

8.6.1. Differences between Copy on Write (COW) and Redirect on Write:

The biggest differences between the two types of snapshots are the underlying technique used to manage the location of the data that needs to be preserved. When an application writes to a disk using COW, the snapshot engine copies the contents of the block(s) being overwritten to a new location (snapshot pool), whereas Redirect on Write snapshots write the new data to a new page in the storage pool and updates the internal pointers to indicate the location of the most up to date copy of the data.

Where COW requires three I/O's to complete the operation, Redirect on Write only requires one I/O (write the new data to disk). The advantage of this is the host application does not suffer the inherent performance penalties that exist with the legacy COW snapshots.

8.6.2. Redirect on write snapshot overview:

A snapshot is in theory a new volume with pointers that reference the original volume's data. To do this we create a new volume for all new incoming writes, create page tables and data structures within the storage pool that duplicate the original volume to keep track of all references for the volume.

A reference count in the page tables is used for each allocated page in a volume to record how many volumes (snapshots) use it. Since reference counts are shared by multiple volumes (snapshots), they are not stored with the volume. They are instead kept in the component object in cache for fast access.

If the volume reference count is 1 then writes will continue as if the volume were a standard volume. If the reference count increases copy services will know a snapshot of that volume has been taken and allocate a new page from the storage pool. When a snapshot is deleted the volume reference count is reduced by 1 and the relevant volume (snapshot) is deleted.

Some Frequently Asked Questions about the new design:

Q: Is a snapshot pool required to use the new copy service?

A: No, the new design removes the requirement for a snapshot pool. There is a small volume created within the storage pool for pointers.

Q: Are there any performance penalties using the new copy services?

- A: No, as each volume (or snapshot) is treated as a standard volume with reference pointers back to the source volume, snapshots have the same performance profile as the original source volume.
- Q: Are there performance implications when deleting one or multiple snapshots?
- A: No, all information relating to snapshots is using reference pointers, once these pointers are deleted the virtual storage services de-allocates the required pages from the overall system map.
- Q: Can a snapshot of a snapshot be taken?
- A: Yes, due to the new design of copy services each snapshot is treated as a peer volume and has all the functionality of the original volume.
- Q: Can the new copy services snapshots be used with standard linear VDisks?
- A: No, virtual storage and storage pools are not compatible with traditional VDisks
- Q: Can multiple volumes be snapshotted at the same time?
- A: Yes, this functionally is similar the existing DMS functions.
- Q: Are snapshots with the new copy services application consistent?
- A: Yes, new copy services support Microsoft VSS and VMware VAAI, consistency outside of these areas will require scripting and sample scripts can be provided.

8.7. RealCache: SSD Read Cache

8.7.1. Read Cache

A read cache is a special-purpose disk group dedicated to hold volatile copies of data within one or more SSDs. An SSD read cache is a new concept in RealStor.

Some advantages of a read cache are:

- Improved performance in read-centric workloads where maximum IOPS is crucial
- Requires fewer SSDs when compared to a dedicated SSD tier.
- A single SSD disk can be used as a read cache for a storage pool; whereas, an SSD tier requires multiple disk to provide fault tolerance.
- Takes advantage of tiering-like algorithms so that only active pages are copied into the read cache.
- Active pages can be from any volume within the pool.

The read cache becomes active when the controller cache is saturated. The controller cache software gives “hints” to the system software about which pages it would like to keep, but can’t. The system then copies these pages into the read cache. Reads are served from the SSD, and writes are written to both SSD and HDD.

Key concept: The read cache keeps copies of data that already exist in the hard disk tier; therefore, it doesn't require more SSD disks to be fault-tolerant. The SSD disk(s) provide the performance and the HDD disks provide the fault tolerance. This is a cost-effective method of providing performance.

A read cache is created using the Add Disk Group dialog box. From the Storage Management Console:

- Click on the “Pools” icon in the navigation menu to the left.
- Click the “Action” menu, and select the “Add Disk Group” menu item.

8.7.2. Key CLI Commands for Read Cache:

The CLI may be used to create disks groups for Autonomic Tiers (HDD) and Read Cache. Figures 4-6 demonstrate the creation of a read cache, the viewing of disk groups, and the deletion of a disk group respectively. The examples highlight a read cache, but the commands may be used for any type of tier.

Figures 4-6

Creating a disk group

```
# add disk-group type read-cache pool a disks 1.21-22 MyReadCache
```

Viewing disk groups

```
# show disk-groups
```

Removing a disk group

```
# remove disk-groups MyReadCache
```

9. Warranty

9.1. General Warranty Options

9.1.1. 3 year warranty from date of shipment on all enclosures

- 3 Year 7x24 Help Desk Support
Global support providing access to trained support specialist any time you need help with technical or administrative issues
- Access to the SANsolve CRM
Around the clock web services allowing users to submit technical support requests, submit RMA requests, or check status and access information in the knowledge base
- Access to Dot Hill Technical Knowledge base
24 hour web service allowing users to search for technical solutions to known problems. Includes general announcements, firmware/software releases, white papers, technical bulletins and product documentation.
- Advance Exchange
If an equipment problem is confirmed by Dot Hill trained technical support personnel, required replacement part(s) will be shipped in advance of the return of the suspect part(s)

9.2. Support Contact #'s / Resources

9.2.1. Customer Resource Center (CRC)

Dot Hill's Customer Resource Center provides you with easy access to the product documentation you are currently accustomed to finding on the Support section of the Dot Hill web site. Enter the new Customer Resource Center here, or bookmark crc.dothill.com for direct access to product documentation and support materials. Available on the Customer Resource Center are product manuals, installation guides, software downloads, certification information, interoperability matrices, and frequently viewed articles.

Contact Technical Support: support@dothill.com
Phone: +1-877-DOT-7X24 (+1-877-368-7924)
If calling from outside the U.S.A.,
Phone: 001-303-845-3200 – select #1 for support

9.2.2. SANsolve Contact Information

SANSolveHelp@dothill.com or 1-877-368-7924

9.2.3. Open a Support Ticket

<http://www.dothill.com/support/sansolve/>

9.2.4. Submitting RMA Requests

- RMA requests are submitted via SANsolve, Dot Hill's web-based call center interface.
- Users may register for access to SANsolve at <http://www.dothill.com/support/sansolve/new-user-registration>
- Registered users may open an RMA request in SANsolve by selecting Current User Login at <http://www.dothill.com/support/sansolve>
- SANsolve will require the user to provide the part and serial number of suspect part(s), a description of the perceived failure, and the ship to company name, address, contact name and phone number of the recipient of replacement parts

10. Interoperability & Agency Certifications

10.1. Interoperability Certifications:

For the latest Dot Hill Interop Matrix of Hypervisor, Operating System, HBA, and SAN switch certifications, go to <http://CRC.Dothill.com>.

10.1.1. VMware Certifications

<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=san>

10.1.2. Microsoft Certifications

<http://www.windowsservercatalog.com/results.aspx?bCatID=1338&cpID=0&avc=10&OR=1>

10.2. Agency Certifications – Safety, Emissions, Harmonics, Country Approvals:

The following table list certifications received, or in process by chassis.

	2U12	2U24
Safety		
NRTL Cert and Report to UL 60950-1 (USA Safety)	Yes	Yes
NRTL Cert and Report to CSA 22.2 No. 60950-1 2nd Edition (Canadian,	Yes	Yes
NRTL Cert and Report to EN 60950-1 (European Union)	Yes	Yes
NRTL Cert and Report to IEC 60950-1 (International)	Yes	Yes
NRTL Cert and Report to EN 60950-1, Am 1 (GS Mark,Germany)	Yes	Yes
EMC (Emissions, Harmonics, Flicker, Immunity)		
Emissions FCC Class A Mark to CFR 47 Part 15 Subpart B Class A (USA, EMC)	Yes	Yes
Emissions ICES-003 – Class A (Canada)	Yes	Yes
Emissions EN 55022 Class A (EU)	Yes	Yes
Emissions EN 300 386 Class A (EU Telco)	Yes	Yes
Emissions AS/NZS CISPR 22 Class A (Australia, New Zealand)	Yes	Yes
Emissions VCCI Class A (Japan)	Yes	Yes
Emissions GOST R 51318.22 Class A (Russia)	Yes	Yes
Emissions KN 22 Class A (S. Korea)	Yes	Yes
Emissions CNS 13438 Class A (Taiwan)	Yes	Yes
Harmonics EN61000-3-2 (EU)	Yes	Yes
Flicker EN 61000-3-3 (EU)	Yes	Yes
Immunity EN 55024 (EU)	Yes	Yes
Immunity EN 300 386 (EUTelco)	Yes	Yes
Immunity GOST R 51318.24 (Russia)	Yes	Yes
Immunity KN 24 (S. Korea)	Yes	Yes
Other		
RoHS 6/6 Compliance (with exemptions)	Yes	Yes
China RoHS	Yes	Yes
WEEE	Yes	Yes
Country Approvals		
USA	Yes	Yes
Canada	Yes	Yes
EU	Yes	Yes
Australia/New Zealand	Yes	Yes
Japan	Yes	Yes
China	PSU only	PSU only
Russia	EAC Pending	EAC Pending
Mexico	Yes	Yes

Germany	Yes	Yes
South Korea	Pending	Pending
Taiwan	Yes	Yes
India	Yes	Yes

11. Data Sheet Specifications – RBOD (Power, Temp, Acoustic, Dimensions)

RBOD	2U24	2U12
Model #	D3824, D3524	D3834, D3534
Power Requirements - AC Input:		
Input Power Requirements	100-240VAC 50/60Hz	100-240VAC 50/60Hz
Max Input Power	400W maximum continuous	375W maximum continuous
Heat Dissipation	1488 BTUs/hour	1488 BTUs/hour
Power Supply:	Bronze Rated - high efficiency	Bronze Rated - high efficiency
	82% @ 20% load	82% @ 20% load
	86% @ 80% load	86% @ 80% load
	85% @ 100% load	85% @ 100% load
Power Requirements - DC Input:		
Input Power Requirements	-39 to -72VDC, -48/-60V nominal	-39 to -72VDC, -48/-60V nominal
Max Input Power	500W maximum continuous	500W maximum continuous
Heat Dissipation	1706 BTU/hour	1706 BTU/hour
Temperature and Humidity Ranges		
Operating Temperature	41°F to 104°F (5°C to 40°C)	41°F to 104°F (5°C to 40°C)
Non Operational Temperature		
Shipping Temperature	-40°F to 158°F (-40°C to 70°C)	-40°F to 158°F (-40°C to 70°C)
	Note: De-rate 2°C for every km, up to 3000 meters	Note: De-rate 2°C for every km, up to 3000 meters
Operating Humidity	10% to 90% RH @ 104°F (40°C), non-condensing	10% to 90% RH @ 104°F (40°C), non-condensing
Non-Operating Humidity	Up to 93% RH @ 104°F (40°C), non-condensing	Up to 93% RH @ 104°F (40°C), non-condensing
DECLARED ACOUSTIC NOISE LEVELS		
Sound Power	LWAd=6,75 B	LWAd=6,75 B
Sound Pressure	LpAm - 55dBA	LpAm - 55dBA
SHOCK AND VIBRATION		
Shock, Operational	3G's for 10 ms, half sine	3G's for 10 ms, half sine
Shock, Non-operational	10G 11ms, half sine	10G 11ms, half sine
Vibration, Operational	5Hz to 500Hz, 0.21Grms fl at spectrum	5Hz to 500Hz, 0.21Grms fl at spectrum
vibration, Non-Operational	3-365-3Hz, 1.22 Grms, z-axis, 0.85 Grms, X&Y axis shaped spectrum	3-365-3Hz, 1.22 Grms, z-axis, 0.85 Grms, X&Y axis shaped spectrum
Physical:		
# of drives	With twenty four (24) 2.5" Drives	With twelve (12) 3.5" Drives
Depth (excluding cables)	20.46 Inches / 51.9 cm	20.57 Inches / 52.25 cm
Height	3.5 Inches / 8.9 cm	3.43 Inches / 8.71 cm
Width	17.6 Inches / 44.7 cm	18.99 Inches / 48.24 cm
Chassis weight	36.4 lb. / 16.5 kg	58.2 lb. / 26.9 kg
Chassis weight w/Drives	51.8 lb. / 23.5 kg	78.2 lb. / 35.5 kg

12. Data Sheet Specifications – JBOD (Power, Temp, Acoustic, Dimensions)

JBOD	2U24	2U12
Model #	J6G24	J6G12
Power Requirements - AC Input:		
Input Power Requirements	100-240VAC 50/60Hz	100-240VAC 50/60Hz
Max Input Power	436W maximum continuous	436W maximum continuous
Heat Dissipation	1488 BTUs/hour	1488 BTUs/hour
Power Supply:	Bronze Rated - high efficiency	Bronze Rated - high efficiency
	82% @ 20% load	82% @ 20% load
	86% @ 80% load	86% @ 80% load
	85% @ 100% load	85% @ 100% load
Power Requirements - DC Input:		
Input Power Requirements	-39 to -72VDC, -48/-60V nominal	-39 to -72VDC, -48/-60V nominal
Max Input Power	500W maximum continuous	500W maximum continuous
Heat Dissipation	1706 BTU/hour	1706 BTU/hour
Temperature and Humidity Ranges		
Operating Temperature	41°F to 104°F (5°C to 40°C)	41°F to 104°F (5°C to 40°C)
Non Operational Temperature		
Shipping Temperature	-40°F to 158°F (-40°C to 70°C)	-40°F to 158°F (-40°C to 70°C)
	Note: De-rate 2°C for every km, up to 3000 meters	Note: De-rate 2°C for every km, up to 3000 meters
Operating Humidity	10% to 90% RH @ 104°F (40°C), non-condensing	10% to 90% RH @ 104°F (40°C), non-condensing
Non-Operating Humidity	Up to 93% RH @ 104°F (40°C), non-condensing	Up to 93% RH @ 104°F (40°C), non-condensing
DECLARED ACOUSTIC NOISE LEVELS		
Sound Power	LWAd=6,75 B	LWAd=6,75 B
Sound Pressure	LpAm - 55dBA	LpAm - 55dBA
SHOCK AND VIBRATION		
Shock, Operational	3G's for 10 ms, half sine	3G's for 10 ms, half sine
Shock, Non-operational	10G 11ms, half sine	10G 11ms, half sine
Vibration, Operational	5Hz to 500Hz, 0.21Grms fl at spectrum	5Hz to 500Hz, 0.21Grms fl at spectrum
vibration, Non-Operational	3-365-3Hz, 1.22 Grms, z-axis, 0.85 Grms, X&Y axis shaped spectrum	3-365-3Hz, 1.22 Grms, z-axis, 0.85 Grms, X&Y axis shaped spectrum
Physical:		
# of drives	With twenty four (24) 2.5" Drives	With twelve (12) 3.5" Drives
Depth (excluding cables)	20.46 Inches / 51.9 cm	20.57 Inches / 52.25 cm
Height	3.5 Inches / 8.9 cm	3.43 Inches / 8.71 cm
Width	17.6 Inches / 44.7 cm	18.99 Inches / 48.24 cm
Chassis weight	36.4 lb. / 16.5 kg	58.2 lb. / 26.9 kg
Chassis weight w/Drives	51.8 lb. / 23.5 kg	78.2 lb. / 35.5 kg

13. General

13.1. *NEBS level-3 certification*

Dot Hill uses UL for Safety Certifications and Nationally Recognized Test Lab (NRTL) for EMC and NEBS compliance testing. Dot Hill 2U12 and 2U24 products also have TUV-GS mark for ease of shipment within Germany. Dot Hill uses the following definitions for products regarding NEBS:

Dot Hill Systems (DHS) Data Storage Arrays may be designed to meet NEBS requirements depending upon customer demand. In order to simplify the definitions of how Dot Hill documents different NEBS requirements, the definitions used below document what this means for Dot Hill products:

- 1) Compliant
 - i. indicates we have designed the product to pass the standard but have not completed all the necessary testing.
- 2) Declaration of Conformance
 - i. indicates that all necessary testing is complete and DHS “DECLARES” we have conducted all testing to the standards.
- 3) Certified NEBS
 - i. indicates that all necessary nesting is complete and DHS “DECLARES” we have conducted all testing to the standards in a NEBS certified test laboratory.

Dot Hill 2U12 and 2U24 product lines are NEBS Compliant.

13.2. *Super Capacitors*

Dot Hills **EcoStor™** Battery-Free Protection provides battery elimination with patented flash alternative (aka super capacitors) providing 10+ years in service time compared to 2-3 year battery life. There is a single Super Capacitor within each controller from Dot Hill. Additional benefits for Super Capacitors include:

- Reduces costs:
- No downtime for replacement
- No FRU to track & stock
- No hazardous waste disposal
- No battery replacement cost
- Elimination of service calls

Super Capacitors are a competitive advantage for Dot Hill. Some competitors use battery packs and its 10Wh versus Dot Hill's 0.1Wh. In other words, competitors need 50-100x the energy to back up write cache to NV memory. They also need 50-100x the time to recharge a fully dissipated battery before they can be protecting the write cache after a power on.

13.3. *SimulCache™ Technology*

SimulCache leverages redundant RAID controllers, thus eliminating the performance degradation associated with conventional cache mirroring. The implementation of a high-speed dedicated PCI-Express bus streamlines write performance resulting in superior overall system responsiveness driving faster application throughput in database-related deployments. This can benefit companies by increasing the number of transactions they can process - i.e. by booking more airline reservations, or more stock trades -- driving increased revenues. SimulCache™ Technology provides:

Patented High Performance /Low Latency Cache Mirroring

- Better Write Performance
- Mirrored writes 92% as efficient as non-mirrored
- Reduces load on controller processor & I/O Bus
- Better Data Protection with redundancy

13.4. *Other Sources of Information*

AssuredSANTM 3xx4 Series

-  [83-00004787-11-01 VSS VDS.pdf](#)
-  [AssuredSAN 3004 Series Getting Started \(GL105\)](#)
-  [AssuredSAN 3004 Series Getting Started \(GL200\)](#)
-  [AssuredSAN 3004 Series FRU Installation and Replacement Guide \(GL200\)](#)
-  [AssuredSAN 3004 Series FRU Installation and Replacement Guide \(GL105\)](#)
-  [AssuredSAN 3004 Series CLI Reference Guide \(GL200\)](#)
-  [AssuredSAN 3004 Series CLI Reference Guide \(GL105\)](#)
-  [AssuredSAN 3004 Series Obtaining and Installing a License \(GL200\)](#)
-  [AssuredSAN 3004 Series Obtaining and Installing a License \(GL105\)](#)
-  [AssuredSAN 3004 Series Setup Guide \(GL200\)](#)
-  [AssuredSAN 3004 Series Setup Guide \(GL105\)](#)
-  [AssuredSAN 3004 Series Storage Management Guide \(GL200\)](#)
-  [AssuredSAN 3004 Series RAIDar User Guide \(GL105\)](#)
-  [Event Descriptions Reference Guide \(GL200\)](#)
-  [Event Descriptions Reference Guide \(GL105\)](#)