



BrickStor Configuration Guide

Version 22.2



Terms of Use Applicable to the User Documentation

The legal notices, disclaimers, terms of use, and other information contained herein (the "terms") apply only to the information discussed in this documentation (the "Documentation") and your use of it. These terms do not apply to or govern the use of websites controlled by RackTop Systems, Inc. or its subsidiaries (collectively, "RackTop Systems") or any RackTop Systems-provided products. RackTop Systems products are available for purchase and subject to a separate license agreement and/or terms of use containing very different terms and conditions.

Terms of Use and Copyright and Trademark Notices

The copyright in the Documentation is owned by RackTop Systems and is protected by copyright and other intellectual property laws of the United States and other countries. You may use, print out, save on a retrieval system, and otherwise copy and distribute the Documentation solely for non-commercial use, provided that you (i) do not modify the Documentation in any way and (ii) always include RackTop Systems' copyright, trademark, and other proprietary notices, as well as a link to, or print out of, the full contents of this page and its terms.

No part of the Documentation may be used in a compilation or otherwise incorporated into another work or with or into any other documentation or user manuals, or be used to create derivative works, without the express prior written permission of RackTop Systems. RackTop Systems reserves the right to change the terms at any time, and your continued use of the Documentation shall be deemed an acceptance of those terms.

RackTop Systems, the RackTop Systems logo, BrickStor, CyberConverged, and certain other trademarks and logos are trademarks or registered trademarks of RackTop Systems, Inc. in the United States and other countries. Other company, product and service names may be trademarks or service marks of others.

© 2021 RackTop Systems, Inc. All rights reserved.

Disclaimers

The Documentation and any information available from it may include inaccuracies or typographical errors. RackTop Systems may change the documentation from time to time. RackTop Systems makes no representations or warranties about the accuracy or suitability of any RackTop Systems-controlled website, the Documentation and/or any product information. RackTop Systems-controlled websites, the Documentation and all product information are provided "as is" and RackTop Systems disclaims any and all express and implied warranties, including but not limited to warranties of title and the implied warranties of merchantability and/or fitness for a particular purpose. In no event shall RackTop Systems be liable to you for any direct, indirect, incidental, special, exemplary, punitive, or consequential damages (including but not limited to procurement of substitute goods or services, loss of data, loss of profits, and/or business interruptions), arising out of or in any way related to RackTop Systems-controlled websites or the documentation, no matter how caused and/or whether based on contract, strict liability, negligence or other tortious activity, or any other theory of liability, even if RackTop Systems is advised of the possibility of such damages. because some states/jurisdictions do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitations may not apply to you.

Release Date: 15 January 2021

Table of Contents

Getting Started with BrickStor	5
BrickStor Appliances and Components	5
Open Communication Ports Requirements	7
Initial Configuration	8
Logging into the BrickStor Security Platform using BrickStor SP Manager	8
BrickStor SP Manager	11
General User Layout and Conventions	11
The Rack View Interface	15
General Appliance Information	18
Appliance Sharing Information	19
Appliance Data Protection Information	20
Network Information	21
System Information	22
Self Encrypting Drives	24
Drive Enrollment	24
Other Self Encrypting Drive Operations	25
Exporting and Backing Up Keys	25
Cryptographically Erasing SEDs	26
SED Protection on the Main Pane	27
Pools	28
Boot Pools	28
Hybrid Pools	28
RAID Performance	30
Pool Hierarchy and Containers	32
Pool Types	32
Creating Pools	32
Viewing Pools	35
Managing Pools	36
Pool Storage Utilization	49
Pool Performance	50
Pool Sharing Information	52
Pool Settings	52
Destroying Pools	53
Datasets	55
Shares	55
Creating Datasets	57
Working with Datasets	59
Dataset Permissions	59
Quotas and Reservations	63
Dataset Storage Utilization	64
iSCSI	64
Encryption and Key Management	67

Managing Encryption	67
Encryption Best Practices	68
Snapshots	70
Snapshot Indexing	70
Restoring a File from a Snapshot	70
Snapshot Holds	71
Clones	72
Replication	73
Replication Best Practices	73
Understanding Peers	73
Configuring a Peer Relationship	73
Understanding Peer Status	75
Data Protection Replication	78
Data Protection Policy Configurations	78
Configure the Data Protection Policy for a Storage Profile	79
Managing Replication Details	79
Replication Transfer History	80
Auto Snapshot Data Protection	81
User Behavior Auditing and Analysis	84
Enabling User Behavior	84
User Behavior Audit	85
High Availability	87
High Availability Components	87
HA Cluster Architecture	89
HA Scenarios	90
High Availability (HA) Best Practices	92
Configuring High Availability	93
Prerequisites	93
Setting up Witness Server	93
Forming HA Cluster	95
Managing High Availability	98
HA Cluster Settings	98
Disabling and Enabling HA Head Nodes	100
Managing Resource Groups	102
iSCSI Initiator	110
Configuring iSCSI Initiator	110
Continuing to configure the iSCSI Initiator	111
Adding an iSCSI Target(s)	112
Compliance Reports	115
Accessing Compliance Reports	115
Select Reports by Category	115
Favorite Reports	115
Export Reports	115
Audit Log	116

Accessing the Audit Log	116
Metrics	118
Accessing Metrics	118
Licensing	119
Using the Licensing feature	119
Command Line Reference	120
Configuring Ethernet Address on Physical Interfaces	120
Configuring Default Gateway	122
BSRAPID Configuration	123
Time Zone Setup	123
NTP Setup	123
Hosts Entries	124
RMM (Remote Terminal) IP Address	124
Add Local Accounts to Bsradmins Group	125
Adding and removing e-mail addresses from Report Notification List	125
Joining Active Directory	125
Performance Monitoring	127

Getting Started with BrickStor

The BrickStor Security Platform (BrickStor SP) is a CyberConverged™ network attached storage (NAS) solution that fuses scalable capacity and performance with advanced data security and compliance capabilities. BrickStor eliminates attack vectors present in traditional storage systems while automatically ensuring continuous compliance through storage-based data profiles.

A typical BrickStor deployment consists of:

- At least one controller that provides a centralized management console, via the BrickStor SP Manager user interface, and a database repository for the BrickStor platform.
- At least one enclosure that contains some number of drives for storage capabilities.

In addition to the basic standard configuration, BrickStor can be deployed in a High Availability configuration.

This guide provides information about the features and functionality of the BrickStor Security Platform. The explanatory text, graphics, and procedures in each topic provide detailed information to help you navigate the user interface, maximize the performance of your system, and troubleshoot complications.

The topics that follow introduce you to the BrickStor Security Platform, describe its key components, explain how to log in and out of the system, and help you understand how to use this guide:

- [BrickStor Appliances and Components](#)
- [Open Communication Ports Requirements](#)
- [Initial Configuration](#)
- [Logging into the BrickStor Security Platform using BrickStor SP Manager](#)

BrickStor Appliances and Components

A BrickStor appliance is either a physical server or virtual machine running the BrickStor SP Operating System.

For more information on a hardware BrickStor, see the following topics:

- [Controllers](#)
- [Enclosures](#)
- [Drives](#)

Controllers

The controller contains the BrickStor SP Operating System (BrickStorOS) and provides a centralized management point for your storage deployment and services. Controllers are sometimes referred to as heads, or nodes.

A typical controller is equipped with multi-core Intel CPUs and 256GB or more memory. The system uses this memory for caching, which is discussed in greater detail later in this documentation. Controllers provide networking via onboard interfaces with a typical system containing two 10GbE

Ethernet interfaces onboard, and two or more 10GbE or faster Ethernet interfaces as add-on components for data access. Controllers also provide component redundancy wherever possible, including power, cooling, and storage used by the operating system, etc.

BrickStorOS

BrickStorOS is the Operating System for your BrickStor appliance. It is not a general-purpose operating system. Instead, it serves as part of an embedded system, which in combination with RackTop hardware becomes the BrickStor Security Platform. BrickStorOS provides a console mode, as well as shell access. However, these features exist for supporting very low-level functionality, such as networking configuration, system optimization, troubleshooting, and other diagnostic functions.

WARNING

When attempting to perform actions within the BrickStorOS that are not documented or recommended by RackTop, be aware that these actions may result in system instability, loss of data, and violation of the terms of the system's maintenance contract.

Enclosures

A BrickStor disk enclosure is an appliance with redundant components, which like a controller is engineered to be fault-tolerant. An enclosure is sometimes referred to as a shelf. Enclosures are either fully or partially populated with mechanical and/or solid-state drives. These drives act as the primary storage for your BrickStor Security Platform and are organized into logical groupings called Pools. Enclosures can also contain special cache and write optimized *journal* devices.

Enclosures are attached to controller(s) via dual SAS host controllers, and utilize SAS drives, which permit multi-pathing throughout the system. Multi-pathing adds to system redundancy and IO load distribution. Loss of path to storage may cause a pause, while the system recovers from the loss and continues operating with a single remaining path. Whenever possible, RackTop recommends using multi-pathing throughout your deployment.

Drives

Enclosures are populated with high capacity storage drives. Typical configurations include mechanical Hard Disk Drives, Solid State Drives, or a combination of the two (Hybrid).

In some instances, special purpose drives used for caching or journaling are installed in the controller. These are often referred to as Write Cache or Read Cache.

Both types of drives use SAS interfaces, which possess dual-ported capability and enables multi-pathing as described in [Enclosures](#). Enterprise grade drives are a standard feature in all systems and are selected to fit a specific configuration both in terms of capacity and parity scheme or mirroring.

High Availability

There are high availability options available in addition to the basic standard configuration. High availability is a configuration which includes two controllers and one or more disk enclosures with shared access between these controllers. The basic premise is high availability to some degree protects from catastrophic physical failure, or failure in operating system on a controller. Because

storage is common between the controllers, high availability configuration is not meant to provide increased protection for storage, instead storage is protected through mirroring or a parity scheme such as RAID.

Open Communication Ports Requirements

By default, the following ports are open to allow BrickStor to take advantage of various features and functionality. The following table lists these ports.

Table 1. BrickStor Open Communication Ports Requirements

Ports	Description/Service	Protocol	Direction	This port is open to/Purpose
22	SSH	TCP	inbound	Receive Management and Replication data
22, 8444	TCP Replication	TCP	outbound	Send Replication
25, 587	mail	TCP	outbound	send notification emails
53	DNS	UDP	bidirectional	Domain name Service
88	Kerberos	UDP	outbound	Authentication
111	NFS/rpc	TCP/UDP	inbound	NFS client access
123	NTP	UDP	bidirectional	Time synchronization
139, 445	SMB	TCP/UDP	inbound	SMB/CIFS client access
161	SNMP	UDP	bidirectional	Monitoring with SNMP
162	SNMP traps	UDP	outbound	Sending alerts to SNMP stations
389, 636	LDAP	TCP/UDP	outbound	Access to directory service servers
443	HTTPS	TCP	outbound	Call Home for Software Updates (https://myracktop.com)
514	syslog	TCP/UDP	outbound	Logging
548	AFP	TCP	inbound	Apple client access
2049	NFS/portmap	TCP/UDP	inbound	NFS client access

Ports	Description/Service	Protocol	Direction	This port is open to/Purpose
3205, 3260	iSCSI	TCP	inbound	iSCSI client/initiator access
4045	NFS/lockmgr	TCP/UDP	inbound	NFS client access
4746	hiavd	TCP	bidirectional	High Availability (between HA nodes)
5696, 8445	KMIP	TCP	outbound	Access to key management server
8086, 8088	influxdb	TCP	inbound	Used for BrickStor SP Manager (charts)
8443	bsrapid	TCP	inbound	Used for BrickStor SP Manager (https)

Initial Configuration

On the first boot or when the appliance is not licensed the Out of Box Experience (OOBE) utility will show. It will step you through configuring management network and license the appliance. See [Out of Box Experience](#) page for more details.

Default Login

The default appliance user name is **root** and the password is “**racktop**”. This password is well known and should be changed immediately.

Logging into the BrickStor Security Platform using BrickStor SP Manager


BrickStor has a user interface called BrickStor SP Manager that you can use to perform administrative, management, analysis, and auditing tasks. BrickStor SP Manager can manage a single BrickStor or multiple BrickStors. BrickStor SP Manager runs on Microsoft Windows.

To download and install BrickStor SP Manager, use a web browser and enter the IP address or host name of the appliance. The default web page on the appliance contains downloadable links to the BrickStor SP Manager along with some other resources discussed later in this guide.


The BrickStor SP Manager zip file can be extracted into any folder and will run as a standalone client without an install. The **brickstorspmgr.exe** file in the extracted folder is the executable program.

To log into BrickStor via BrickStor SP Manager:

1. Run the brickstorspmgr.exe by double clicking.



2. Click **Add Connection** to add a connection to a BrickStor SP appliance.



3. In the Add Connection dialog box, enter the following:
- For authentication server, enter the system’s IP address or host name.
 - Enter your username.
 - Enter your password.

- Optionally, select whether to have BrickStor SP Manager save your password for subsequent logins.
4. If you have already connected a BrickStor instance, click **login** for that instance.
 5. In the Connect To dialog box, do the following:
 - Verify the system's IP address.
 - Verify your username.
 - Enter your password.
 - Optionally, select whether to have BrickStor SP Manager save your password for subsequent logins.

BrickStor SP Manager


BrickStor SP Manager provides the user interface for configuring and managing your BrickStor deployment. BrickStor SP Manager is a responsive and context-aware interface that allows you to drill down and manage your BrickStor to a granular level. You can use BrickStor SP Manager to manage a single BrickStor or multiple appliances.

The topics that follow provide a basic interface tour that this guide will build upon in subsequent topics:

- [General User Layout and Conventions](#)
- [The Rack View Interface](#)

General User Layout and Conventions

The BrickStor SP Manager interface is divided into three panes which are described below:



1. the Connections pane
2. the Details pane
3. the Changes pane

Connections Pane

The Connections pane allows you to connect to BrickStor appliances, and navigate their pools and datasets.

Details Pane

The Details pane allows you to configure and manage storage, security, and compliance features.

The tabs and menus available in the Details pane are based on the selection made in the Connections pane. When the top-level Appliance/Node is selected, the system displays different menu tabs than when a pool or dataset is selected for example. Also, certain tabs, such as user behavior, will not be visible if the feature is not enabled. The hierarchy of the Connections and tabs is Appliance, then Pool, and then Dataset. If a menu such as user behavior is selected at the pool level, the system will display all activity related to the pool. However, if you select it at the dataset level, the scope will be narrowed to the dataset. Menus and tabs are relative to position within the interface.

Instead of taking a deep dive into the Details pane here, this documentation covers the tabs and menus herein where it aligns with particular features.

Changes Pane

After you make any configuration changes, they appear in the Changes pane for final review and commit. BrickStor SP Manager does not make actual changes to BrickStor until you commit those changes. Changes that make data unavailable or destroy data require you to acknowledge the possible negative effects before the commit button becomes active.

NOTE

Changes to high availability and resource group movements are not processed through the commit queue.


Main Menu

In the BrickStor SP Manager title bar, you can access the following options:

- [About Menu](#)
- [Search Menu](#)
- [View Menu](#)

About Menu

The About Menu displays BrickStor SP Manager information.

**TIP**

By setting a value, for example 5GB, in the Trace Query and Commit box will create a local log on the machine running BrickStor SP Manager with all of the GUI requests and responses.

Search Menu

The Search Menu allows you to search through your current BrickStor appliance for pools, datasets, etc.

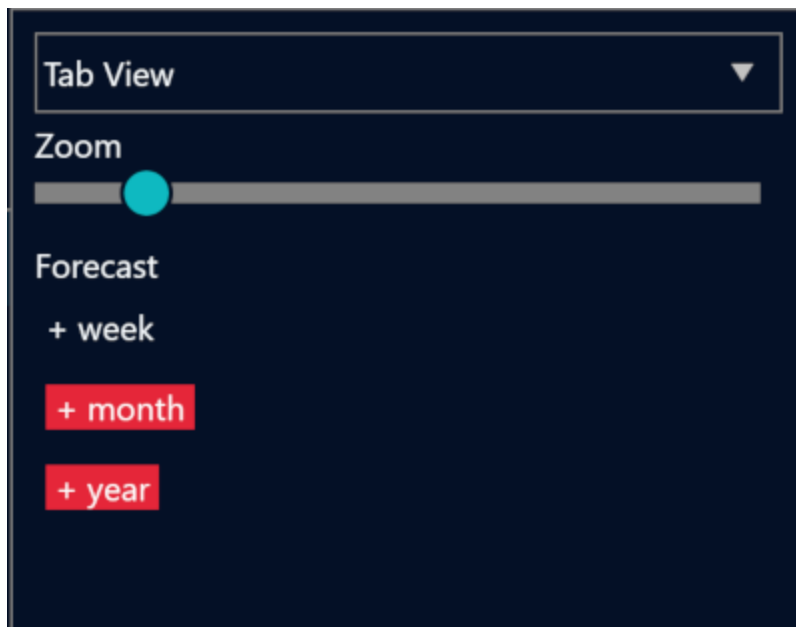
The screenshot displays the BrickStor SP Manager interface. On the left, a sidebar shows the system hierarchy: rts-demo-bar-01 (10.1.19.1) with three pools (p01, p02, p03) and their respective free space. The main area is titled 'Search Results:' and features a search bar, filters (any status, pools, file systems, volumes), and a table of search results. A line graph shows usage over time, and a summary card highlights '01 Dave Archive' with 1.38TB. A 'View Menu' is visible at the bottom of the summary card.

Path	Used	Free
p01	88KB	16TB
p01/global	128KB	15.9TB
p01/global/Josh	1.08GB	15.9TB
p01/replication	96KB	15.9TB
p01/global/frank	296KB	4TB
p01/global/Traudt	30GB	15.9TB
p01/global/DaveDemo	40.8GB	15.9TB
p01/global/ScottDemo	34.3GB	15.9TB
p01/global/jonathanGF	2.38GB	15.9TB
p01/global/Kirby-Demo	390.9MB	15.9TB
p01/global/02 Juno TDM	4.86GB	15.9TB
p01/global/jonathanlab	111.3GB	15.9TB
p01/global/frank/custom	10.2MB	4TB
p01/global/frank/simple	5.41MB	4TB
p01/global/frank/MySpace	280KB	4TB
p01/global/frank/Lab Test	18.3MB	4TB
p01/global/Kirby-TestDemo	365.6MB	15.9TB
p01/global/SecureGFS-demo	324KB	15.9TB
p01/global/01 Dave Archive	1.38TB	15.9TB
p01/global/frank/HIPAA Data	480KB	4TB
p01/global/001 Dave Archive NFS	246.3MB	15.9TB
p02	352KB	262GB
p02/global	864KB	171.1GB
p02/replication	384KB	171.1GB
p02/global/VMTDM	8GB	171.1GB

View Menu

The View Menu allows you to change the BrickStor SP Manager layout. You can choose between Tab View (default) and Flow View, which displays all sections next to each other. You can also view forecast data for the system.

TIP | Tab view is recommended for normal administration on small screens.




Finally, you can adjust Zoom properties, which change the width of columns in all views.

The Rack View Interface

Rack View displays a graphical representation of your current BrickStor hardware, including any controllers, enclosures, and drives that are within these appliances.

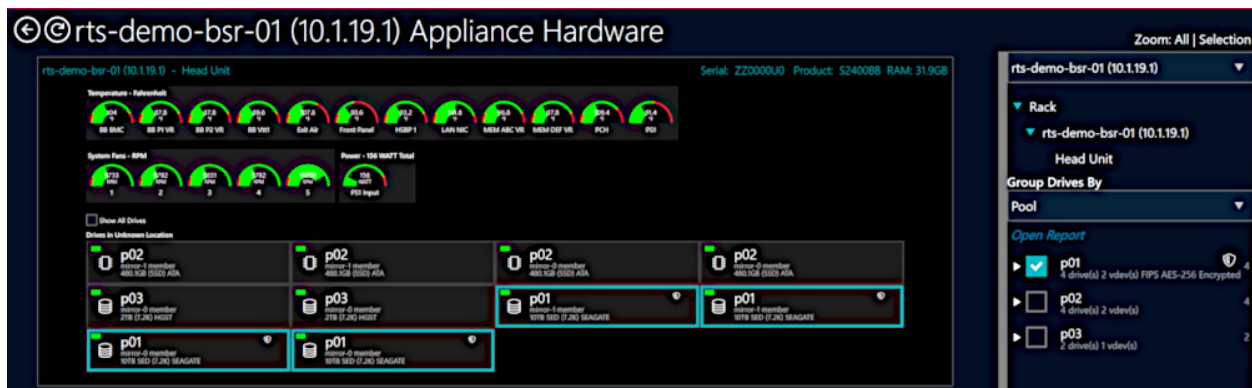
To access Rack View, choose the the appliance in the Connections pane, then click the Rack View link at the bottom of the Details pane.

Accessing Rack View



You can use Rack View to easily view and modify your appliance hardware. Rack View allows users to add or modify pools and vdevs and gives visuals that allow users to see what changes will occur to the system's hardware prior to committing them. It will also display various diagnostic information such as the values of temperature sensors in the system and the fan speeds. On the upper right-hand side, you can select which appliance you want to zoom to. The appliance will be highlighted in yellow when the mouse is hovered over it and left clicking will zoom to the appliance.

The right-hand side of Rack View also allows you to group the drives in the appliances based on certain properties such as pool, make, and vdev type. To change the grouping type, select the dropdown under Group Drives By and then select how you want to group them. When hovering over one of these Drives groups, affiliated drives will be highlighted and left clicking will zoom to the drives. You can also expand these groups with the arrow and select individual drives that are a part of the group.



Accessing Rack View

You can access Rack View from either the Connections or the Details pane.

Accessing Rack View from the Connections pane

To access Rack View from the Connections pane, complete the following steps:

1. From the Connections pane, select either the appliance level or the pool level.
2. Right-click and choose one of the following options:
 - At the appliance level, right-click and select **Open Rack View**.
 - At the pool level, right-click and select **Open Pool Rack View**.

Accessing Rack View from the Details Pane

To access Rack View, complete the following steps:


1. From the Connections pane, select either the appliance level or the pool level.
2. In the lower portion of the details pane, click **Rack View**.

Toggling Identifying Lights


Rack View allows you to toggle a physical indicating light on each drive to assist with identifying the correct drives on the machine. You can either select one drive by clicking directly on it in Rack View, or multiple drives using the Group Drives By interface on the right-hand side. Once the appropriate drives have been selected click the ident on button at the bottom of the screen.



This will bring up the Enable bay indicator LEDs dialog box, where you can turn on the lights for either the selected bays, bays with unknown drives, or bays without drives. You can also choose to disable all other indicator lights to ensure only the desired drives have their lights enabled.



Drives with their indicating LEDs enabled will have a blinking orange indicator on Rack View as well as on the physical drive on the appliance.



To disable the identifying lights, select the desired drives like before and click the ident off button.

This will bring up the Disable bay indicator LEDs dialog box where you can turn off the lights on either the selected bays, bays with unknown drives, bays without drives, or all bays in general.

General Appliance Information

BrickStor SP Manager allows you view all current problems and warnings with the node and its imported pools. From this view you can see which pools are currently imported and exported on the selected BrickStor instance.

Viewing General Appliance Information

To view BrickStor general information, complete the following steps:

1. From the Connections pane, select the appliance level.
2. In the details pane, select the **General** tab.

The screenshot displays the BrickStor SP Manager interface for an appliance named 'rts-demo-bsr-01 (10.1.19.1)'. The interface is divided into several sections:

- Header:** BrickStor SP CyberConverged™ NAS. The appliance name and IP address are shown at the top right.
- Search:** A search bar is located below the header.
- Left Panel (Connections):** Lists three pools:
 - p01:** 4 drive(s) 2 vdev(s) FIPS AES-256 Encrypted. 16TB free of 17.6TB.
 - p02:** 4 drive(s) 2 vdev(s). 262GB free of 860.2GB.
 - p03:** 2 drive(s) 1 vdev(s). 1.74TB free of 1.76TB.
- Right Panel (Details):**
 - General:** Shows 'Data' with 2.13TB and 'R. Shared - Non-Reserved Pool Free Space' with 17.7TB.
 - Sharing:** 25 SMB shares, 5 NFS shares.
 - Auto Snapshot Data Protection:** (Status icon)
 - Replication:** (Status icon)
 - Encryption:** 4 encrypted drives, 17 encrypted datasets.
 - Metrics:** (Status icon)
 - Audit:** (Status icon)
 - Network:** (Status icon)
 - TDM:** 1 location(s).
 - System:** (Status icon)
- Pools Section:**
 - p01:** 4 drive(s) 2 vdev(s) FIPS AES-256 Encrypted. 16TB free of 17.6TB.
 - p02:** 4 drive(s) 2 vdev(s). 262GB free of 860.2GB.
 - p03:** 2 drive(s) 1 vdev(s). 1.74TB free of 1.76TB.

Appliance Sharing Information


At the appliance level, the Sharing tab allows you to view all shares currently on an appliance by protocol. In addition, you can view if the datasets are encrypted and on self-encrypting drives. This view also provides a status of the protocol services and health.



Viewing Appliance Sharing Information

To view BrickStor Sharing information at the appliance level, complete the following steps:

1. From the Connections pane, select the appliance level.
2. In the details pane, select the **Sharing** tab.



Appliance Data Protection Information


Data protection encompasses snapshots and snapshot replication. From this tab the admin can monitor data protection health and status for the node as well as configure replication and policies. This tab shows the status of data protection services, peers, policies, and recent restorations. On the Data Protection screen, you can:

1. View the status of Data Protection and its services
2. View and drill down into Replication Peers
3. View the current status of Replication Tasks

Viewing Appliance Data Protection Information

To view BrickStor Data Protection information at the appliance level, complete the following steps:

1. From the Connections pane, select the appliance level.
2. In the details pane, select the **Data Protection** tab.




Network Information

BrickStor SP Manager allows you to view all of the interfaces in your BrickStor deployment. A healthy system should display a green status indicator for all vnics. Each interface displays the IP, interface name, physical interface or aggregate where the vnic resides, MTU size, and port speed.

Viewing Network Information

To view BrickStor network information, complete the following steps:

1. From the Connections pane, select the appliance level.
2. In the details pane, select the **Network** tab.



System Information

BrickStor SP Manager allows you to view system information, service status, and the BrickStor operating systems available for download and installation.

On the service tab, you can find your customer ID, Serial Number and the running version of the OS when calling support. From this admin can all power off and reboot the node as well as access compliance reports. It is from this tab that the admin configures the HA Cluster once the command line steps have been completed. See HA Cluster Configuration for cluster setup details.

Viewing System Information

To view BrickStor system information, complete the following steps:

1. From the Connections pane, select the appliance level.
2. In the details pane, select the **System** tab.

General

Sharing
3 SMB shares
3 NFS shares

Auto Snapshot Data Protection

Replication

Encryption

Metrics

Audit

Network

System

Hardware

Customer: CN000001
 Manufacturer: RackTop Systems Inc.
 Product: 440BX Desktop Reference Platform
 OS: BrickStorOS 22.0 Test (Build 14)
 Serial Number: ZZ0000Z6
 Type: Virtual Appliance
 Time Zone: America/New_York

bp (system)
1 drive(s) 1 vdev(s)

38.1GB free of 61.5GB

Setup HA Cluster

Export Encryption Keys

System: Reboot Shutdown

[Rack View](#)

[Compliance Reports](#)

Peers +
None

Allow peer connections from +
None

Services

- System Services ▼
- Encryption Services ▼
- Data Protection Services ▼
- SMB Services ▼
- NFS Services ▼
- AFP Services ▶ ▼
- iSCSI Services ▼
- HA Services ▼

[Rack View](#)

+ ⬇
add pool export pool

OS Version	Status
22.0 Test (Build 14)	Running
22.0 Test (Build 28)	Available
23.0 Test (Build 1395)	Available
23.0 Developer Build (Build 1392)	Available
23.0 Test (Build 1391)	Available
22.0 Test (Build 27)	Available
23.0 Test (Build 1390)	Available
22.0 Test (Build 26)	Available
23.0 Test (Build 1389)	Available
22.0 Test (Build 25)	Available
22.0 Test (Build 24)	Available
23.0 Test (Build 1388)	Available
23.0 Developer Build (Build 1387)	Available
22.0 Test (Build 23)	Available
21.2.1	Available
23.0 Test (Build 1381)	Available
22.0 Test (Build 22)	Available
23.0 Test (Build 1380)	Available
23.0 Test (Build 1379)	Available
23.0 Test (Build 1378)	Available
22.0 Test (Build 20)	Available
23.0 Test (Build 1376)	Available
23.0 Developer Build (Build 1373)	Available
22.0 Test (Build 19)	Available
23.0 Test (Build 1372)	Available
23.0 Developer Build (Build 1369)	Available
23.0 Test (Build 1370)	Available
23.0 Developer Build (Build 1368)	Available
23.0 Test (Build 1367)	Available
22.0 Test (Build 13)	Available
22.0 Test (Build 12)	Available
23.0 Test (Build 1366)	Available
23.0 Test (Build 1365)	Available
22.0 Test (Build 11)	Available

Self Encrypting Drives


BrickStor can leverage TCG FIPS 140-2 certified self-encrypting drives for increased security. To manage the keys and disks within BrickStorOS does require a special license from RackTop and appropriate FIPS drives. TCG licensed systems may come with drives encrypted using a factory generated key. Self-Encrypting Drives placed in a system that are not licensed will not lock when power is removed.

TCG Must be licensed and the Key Manager must be properly configured before you can utilize this feature


BrickStor SP supports local and external key management. See [Encryption and Key Management](#) for more details.

Drive Enrollment

Once the key manager is configured drives can be enrolled in the system. Each drive will receive a unique key used to unlock the self-encrypting drive known as the key encryption key (KEK) from the key manager and configure the drive to auto lock when power is removed from the drive. To enroll drives or a pool in the system go to the hardware view page of the UI. If you select a drive that is not in a pool you can select multiple drives and enroll the ones you choose to enroll. If you select a drive that is already a member of a pool it will enroll all drives that are a member of that pool.



Other Self Encrypting Drive Operations



Unenroll – Removes drive from SED management and sets the drive to default PIN and sets the drive to stay unlocked.

Rekey – Requests a new key from the key manager and changes the KEK PIN on the drive.

Verify Key – Verify the KEK unlocks the drive and is available from the key management service.

Export Keys – Will provide a password protected file with the KEK PINS that can be imported later for backup purposes or to another node so that the other node can unlock the drives. This is required in HA using the internal key management service.

Import Keys – Allows you to import keys that were exported from the same node or another node into the internal key management database. This is performed for HA nodes to share keys between the heads. This can also be used to import keys to a replacement head node.

Exporting and Backing Up Keys

When using the BrickStor internal key manager it is important to back up the keys and store them in an alternate location.

The `/etc/racktop/keymgrd.conf` file allows users to set the location of the internal key file.


The configuration file also allows users to configure the BrickStor to rotate KEKs on a scheduled interval. This is only recommended when using an external key manager in order to ensure you have backup copies of the keys.

Cryptographically Erasing SEDs

Users can Crypto Erase SEDs which will reset the pins and put them in an unenrolled state. To manage the drive again just enroll the drive.

As part of a pool destroy users can select the crypto erase option. This option is irreversible. Data is permanently destroyed and unrecoverable. However, if you don't select the crypto erase option the data is potentially recoverable in the future off each drive.

If the KEK PIN has been lost for a drive a crypto erase is the only option to put the drive back into a usable state because the drive will become erased and unlocked.



SED Protection on the Main Pane

The screenshot displays the 'Key Manager' and 'Drive Encryption (SED)' sections of the BrickStor SP Manager interface. The 'Key Manager' section includes buttons for 'Export All Encryption Keys', 'Import Encryption Keys', and 'Resync Encryption Keys with Peers'. It also shows 'Auto send key material to' and 'Receive key material from' options, both currently set to 'habsr01'. The 'Drive Encryption (SED)' section provides a summary of drive encryption status: 1 enrolled (FIPS AES-256 Encrypted), 1 unlocked (Ready to Auto-Lock), 13 not enrolled, and 6 not supported. It includes links for 'Drive Encryption Report' and 'Drive Status Report', and buttons for 'Enroll', 'Verify Keys', 'Rekey', 'Export SED Keys', 'Unenroll', and 'Config (Advanced)'. The 'Encrypted Pools' section shows 'onedrive' with 1 drive(s) and 1 vdev(s) FIPS AES-256 Encrypted. At the bottom, 'Encryption Services' is shown as active.

Under the general tab of BrickStor SP Manager users can perform various SED configuration options as well review reports about which drives are enrolled in SED management and the current status of each drive.

Pools

Pools organize storage drives into logical groupings for data management. Pools serve as the containers for your datasets in BrickStor.

There are two types of pools in BrickStor:

- Boot Pools
- Hybrid Pools

BrickStor uses the Boot Pool primarily for appliance administration purposes. For the purposes of data management, when this documentation refers to pools, it is referring to hybrid pools.

Boot Pools

The Boot Pool consists of two mirrored SSDs and contains the BrickStorOS. It is a mirrored pool used to boot the appliance. This pool should remain untouched during normal BrickStor operations. Logs stored on the boot pool are set to auto rotate and expire to prevent any partition or directory from becoming full.

Hybrid Pools

A typical BrickStor deployment is referred to as a *hybrid storage system*. A *hybrid pool* is a collection of drives, optionally with dedicated read-optimized cache devices and write optimized journal devices. All storage pools are hybrid pools because they are a combination of in-memory read cache as well as actual high capacity persistent storage and optionally read and write cache devices. The high capacity data drives are organized into virtual devices called vdevs.

A vdev, also known as a stripe, is a virtual device that can be a single disk, two or more disks that are mirrored, or a group of disks with a parity scheme such as RAID-5. The concept of a vdev is something that abstracts away some unit of storage, which may or may not have any redundancy. vdevs can be viewed as a building block for pools.

Pools are groups of virtual devices usually implemented with some data protection scheme, such as RAID or mirroring, on top of which filesystems and raw block devices are provisioned. A typical hybrid pool is a mix of mechanical drives and solid-state drives. In such a pool, data is redundantly stored on large capacity, slower, typically mechanical devices, arranged into a parity scheme that satisfies data protection as well as capacity and IOPS requirements, while high bandwidth, low latency solid state drives are used for the purposes of caching to accelerate reads and for the purposes of handling synchronous writes, enabling a much better cost to performance ratio over traditional purely mechanical, or purely solid state configurations. BrickStor also configures all flash pools, which continue to leverage RAM for cache solid state disks instead of mechanical disks to provide consistently lower latency and higher IOPS.

You must configure one or more data pools on a system in order to present storage to consumers via AFP, NFS, SMB, etc. While there is no hard limit on number of pools a system can have, usually fewer than four pools are configured on any given system. Under normal circumstances, the burden of designing and configuring pools is not on the customer, but in the instances where a system is no longer satisfying previously prescribed requirements, RackTop strongly recommends that customer contacts support before any changes are made to configuration of any pool.

From a systems administrator's point of view, a pool is a logical organization of independent drives and contains all information about the devices comprising it, including structure, filesystems, raw volumes, replication target if any, etc. This information is encoded within its metadata, which makes it possible to easily migrate pools between systems. Critically, this property means that loss of the controller does not in any way compromise data. A replacement controller is all that's necessary to return to normal operations. This feature also enables BrickStor's high availability capabilities, which can move pools, as well as related network configuration, between nodes in the cluster.

Adaptive Replacement Cache

Adaptive Replacement Cache (ARC) is a portion of memory in the controller dedicated to caching recently accessed data. The ARC caches both recently written data, with the assumption that this data may be read soon after being written as well as recently read data, with the assumption that this data is potentially going to be read again. Depending on the popularity of data it may remain in the cache for a long time, or be evicted in favor of other data, based on criteria which both the user as well as the system can optimize for.

Read Cache

Optional SSD Cache device that can be used to extend the amount of data that is cached for Read operations. When data is evicted from the ARC it will potentially move to the L2ARC (based up on user configuration settings). Data read from L2ARC will be moved back into ARC.

Write Cache

RackTop uses a journal methodology for its write cache and is implemented in most systems as a mirrored SSD pair. A journal is both a software concept and a core physical component, a write ahead log that is used to reduce latency on storage when synchronous writes are issued by clients. RackTop frequently refers to journal as a ZIL, an intent log or a log device. In synchronous write cases, writes are committed to this journal and periodically pushed to primary storage. Journal guarantees that data is protected from loss on power failure due to being in cache before cache is flushed to stable storage.

A log device is normally only ever written to and never read from. A log device i.e. journal is present to protect the system from unexpected interruptions, such as power loss, a system crash, loss of storage connectivity, etc. In rare instances where recovery is necessary due to power loss or some other catastrophe, journal is read from in order to recreate a consistent state of the pool, which may require rolling back some transactions, but results in restoring the pool to a consistent state, unlike traditional storage systems where only best effort is promised. RackTop recommends mirroring journal devices as a means of preventing loss of a journal device, which has performance and potential availability impact. In all pools configured at the factory prior to system shipping, the journal, if present, will be mirrored.

Resilvering

Resilvering is the process of rebuilding a disk within a vdev after a drive has been replaced. BrickStor OS does not have an fsck repair tool equivalent, common on Unix filesystems. Instead, the filesystem has a repair tool called "scrub" which examines and repairs silent corruption and other problems. Scrub can run while the volume is online; scrub checks everything, including metadata and the data. This process works from the top down and only writes data to the disk that is needed. If a disk was temporarily offline it would only have to rebuild the data that was missed while the

device was offline.

RAID Performance

BrickStor uses mirrors and RAID-Z for disk level redundancy within vdevs.

RAIDZ

RAID-Z vdevs are a variant of RAID-5 and RAID-6:

- You can choose the number of data disks and the number of parity disks. Today, the number of parity disks is limited to 3 (RAID-Z3).
- Each data block that is handed over to ZFS is split up into its own stripe of multiple disk blocks at the disk level, across the RAID-Z vdev. This is important to keep in mind: Each individual I/O operation at the file system level will be mapped to multiple, parallel and smaller I/O operations across members of the RAID-Z vdev.
- When writing to a RAID-Z vdev, ZFS will use a best fit algorithm when the vdev is less than 90% full.
- Write transactions in ZFS are always atomic, even when using RAID-Z: Each write operation is only finished if the überblock has been successfully written to disk. This means there's no possibility to suffer from the traditional RAID-5 write hole, in which a power-failure can cause a partially (and therefore broken) written RAID-5 set of blocks.
- Due to the copy-on-write nature of ZFS, there's no read-modify-write cycle for changing blocks on disk: ZFS writes are always full stripe writes to free blocks. This allows ZFS to choose blocks that are in sequence on the disk, essentially turning random writes into sequential writes, maximizing disk write capabilities.

Just like traditional RAID-5 and RAID-6, you can lose up to 1 disk or 2 disks respectively without losing any data using RAID-Z1 and RAID-Z2. And just like ZFS mirroring, for each block at the file system level, ZFS can try to reconstruct data out of partially working disks, as long as it can find a critical number of blocks to reconstruct the original RAID-Z group.

Performance of RAIDZ

When the system writes to a pool it writes to the vdevs in a stripe. A Vdev in a RAID-Z configuration will have the IOPS and performance characteristics of the single slowest disk in that vdev (it will not be a summation of the disks). This is because a read from disk requires a piece of data from every disk in the vdev to complete the read. So, a pool with 3 vdevs in a RAID-Z1 with 5 disks per vDEV will have the raw IOPS performance of 3 disks. You may see better performance than this through caching, but this is the most amount of raw IOPS the pool can deliver from disk. The more vdev's in the pool the better the performance.

Performance of Mirrors

When the vdev's are configured as mirrors the configuration of the pool is equivalent to RAID-10. A pool with mirrored vdev's will always outperform other configurations. A read from disk only needs data from one disk in the mirror. As with RAID-Z, the more vdevs the better performance will be. Resilver times with mirrored vdevs will be faster than with RAID-Z and will have less of a performance impact on the overall system during resilvering. RackTop recommends the use of

mirrored vdevs in environments with high random IO such as virtualization because it provides the highest performance.

Compression

Compression is performed inline and at the block level. It is transparent to all other layers of the storage system. Each block is compressed independently and all-zero blocks are converted into file holes. To prevent “inflation” of already-compressed or incompressible blocks, BrickStor maintains a 12.5% compression ratio threshold below which blocks are written in uncompressed format. BrickStor supports compression via the LZJB, GZIP (levels 1-9), LZE, and LZ4. RackTop finds that LZ4 works very well, balancing speed and compression performance. It is common to realize a 1.3 to 1.6 compression ration with highly compressible data which not only optimizes storage density but also improves write performance due to the reduction in disk IO. RackTop recommends always using compression because any CPU penalty is typically outweighed by the savings in storage and bandwidth to the disk.

Deduplication

Deduplication is performed inline and at the block level, also like compression, deduplication is transparent to all other layers of the storage system. For deduplication to work as expected the blocks written to the system must be aligned. Deduplication even when turned off will not reverse the deduplication of blocks already written to the system. This can only be accomplished through copying or moving the data. Deduplication negatively impacts the system performance if data is not significantly duplicative because an extra operation must be done to look if it is a duplicate block for writes and if it is the last block for deletes. Additionally, the deduplication table must be stored in RAM. This takes up space that could otherwise be used for metadata and caching. Should the deduplication not all fit in RAM then system performance will degrade sharply because every read and write operation will require the system to reread the dedup table from disk.

NOTE | Deduplication is only supported on All Flash Pools.

Clones

ZFS clones create an active version of a snapshot. By creating a snapshot of a base VM and using clones of that same snapshot you can have an unlimited number of copies of the same base virtual machine without taking up more storage capacity. The only increased storage footprint will come from the deltas or differences between clones. Additionally, since each VM will reference the same set of base data blocks the system and user will benefit from caching since all VM’s will be utilizing the same blocks of data.

Imbalance of vdev Capacity

If you wish to grow the capacity of a volume by adding another vdev you should do so by adding a vdev of equivalent size to the other vdevs in the pool. If the other vdevs are already past 90% capacity they will still be slow because data will not automatically balance or spread across all vdevs after the additional capacity is added. To force a rebalance in a VMware environment you can perform a vmotion or storage migration. With the Copy On Write Characteristics of ZFS, the pool will automatically rebalance across all vdevs.

Pool Hierarchy and Containers

Pools include special containers that are used for organizing datasets and volumes so that they always reside within the same location within the pool.

1. Global – Contains all the datasets and other containers except for the tenant containers on a Pool
2. Volume Container – Contains all virtual block devices which are special datasets exposed over iSCSI
3. Replication – Top level container for all incoming replication streams from other pools within the same BrickStor or other BrickStor's
4. Meta – Contains all of the user behavior audit data and the snapshot index data

Pool Types

This in software implementation allows for various parity schemes as well as mirroring configurations. The following are schemes currently supported by RackTop:

The following table explains the pool types that are available in BrickStor:

Table 2. Pool Types

Type	Description
mirror	Equivalent to RAID 10 / RAID 1+0, aka a stripe of mirrors, where two or more drives in a mirror are possible, offers highest availability with a capacity trade-off
raidz3	(triple parity) Like RAIDZ2, but with even more parity protection, allowing for loss of three drives in each group (vdev)
raidz2	(double parity) Equivalent to RAID 60 / RAID 6+0, which allows for loss of two drives in each group (vdev)
raidz1	(single parity) Equivalent to RAID 50 / RAID 5+0, which allows for loss of a single drive in each group (vdev)
disk	(no parity) fast, but with only minimal protection, and total loss if any single device is lost, useful for scratch-only data

Creating Pools

You can create pools from the details pane or Rack View.

Creating Pools from the Details Pane

To create a pool from the details pane, complete the following steps:

1. In Connections, select the appliance.

NOTE On a clean install, only the appliance level will display.

2. In the lower portion of the details pane, click **Add Pool**.

TIP You can also select the General tab, and then click the add icon next to Pools.

The Create Pool dialog box appears.

Create Pool

Name required

Type

mirror

Auto choose drives from alternating enclosures

Drive Type

107.4GB virtual VMWARE

- 1 + vdevs

- 2 + drives per vdev

- 0 + spare drives

Pool name required.

Create Cancel

3. In the Create Pool dialog box, type a name for the pool.
4. Under **Type**, choose one of the following options:
 - mirror
 - raidz3

- raidz2
 - raidz1
 - disk
5. Optionally, select to **Auto choose drives from alternating enclosures** if you want BrickStor SP Manager to select the drives where your pools will reside.

Clear the check box if you prefer to manually select your disks.
 6. Under **Drive Type**, select from available drive types in your deployment.
 7. Select the number of **vdevs**.
 8. Select the number of **drives per vdev**.
 9. Optionally, select the number of **spare drives**.
 10. Click **Create**.
 11. In the Changes pane, click **Commit Changes**.

Creating Pools from Rack View

When you create a pool from Rack View, you can first view a topography of your storage system and then choose drives based on availability.

1. In Connections, select the appliance.

NOTE On a clean install, only the appliance level will appear.

2. Right-click and select **Open Rack View**.
3. In the details pane, select the drives where you want to create a pool.

TIP Shift-click to select multiple drives.

TIP Optionally, selecting a drive from the right-hand dropdown of **Available** when sorted by Pool.

The selected drive will display a blue border.

4. In the lower portion of the Details pane, click **Create Pool**.

The Create Pool dialog box appears.

5. In the Create Pool dialog box, type a name for the pool.
6. Under **Type**, choose one of the following options:
 - mirror
 - raidz3
 - raidz2
 - raidz1
 - disk
7. Optionally, select to **Auto choose drives from alternating enclosures** if you want BrickStor SP Manager to select the drives where your pools will reside.

Uncheck the check box if you prefer to manually select your disks.

8. Under **Drive Type**, select from available drives.
9. Select the number of **vdevs**.
10. Select the number of **drives per vdev**.
11. Select the number of **spare drives** you want the pool to have.
12. Click **Create**.

Rack View will display the queued changes and any pool that will be affected by changes will have the [changes staged] indicator on it.

13. In the Changes pane, click **Commit Changes**.

Viewing Pools

Selecting a pool in the Connections pane displays information about the Pool's structure and performance.

The screenshot displays the BrickStor SP Manager interface for a storage pool named 'p02' on the host 'rts-demo-bsr-01 (10.1.19.1)'. The interface is divided into several sections:

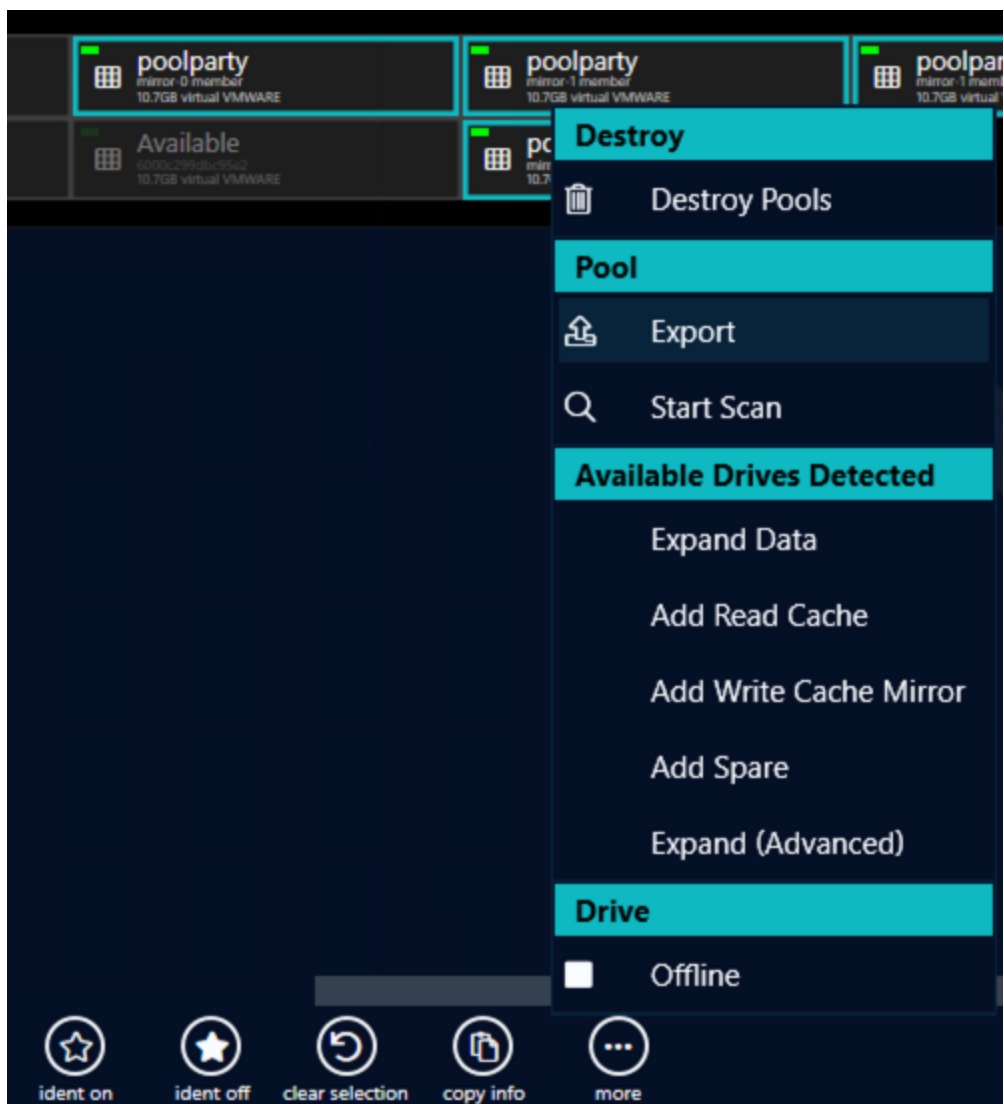
- Header:** Shows the pool name and IP address, along with summary statistics for Data (215.9GB), Reservations (90.9GB), Snapshots (59.6GB), and Shared - Non-Reserved Pool Free Space (493.8GB).
- Left Sidebar:** A navigation pane showing a tree view of storage pools. The selected pool 'p02' is highlighted, showing its free space (96.4GB free of 96.4GB) and its constituent drives.
- Main Content Area:**
 - General:** Displays pool details such as 'Pool' (4 drives, 2 vdevs), 'User Behavior', 'Sharing' (7 descendant SMB shares, 3 descendant NFS shares), 'Auto Snapshot Data Protection' (Enabled), 'Replication' (Disabled), and 'Settings'.
 - Storage Profile:** Shows the pool's configuration, including its location ('rts-demo-bsr-01 (10.1.19.1)') and storage utilization (584.7GB free of 860.2GB).
 - Location:** Lists the pool's location and its constituent drives (p02, 4 drives, 2 vdevs) with their free space (584.7GB free of 860.2GB).
 - Children:** A list of child pools with checkboxes and their free space: 'global' (493.8GB free of 768.2GB), 'meta (system)' (493.8GB free of 494.9GB), and 'replication' (493.8GB free data of 493.8GB).
 - Bottom Bar:** Contains icons for 'snapshot', 'add', 'delete', 'export pool', 'copy into', and 'permissions'.

Managing Pools

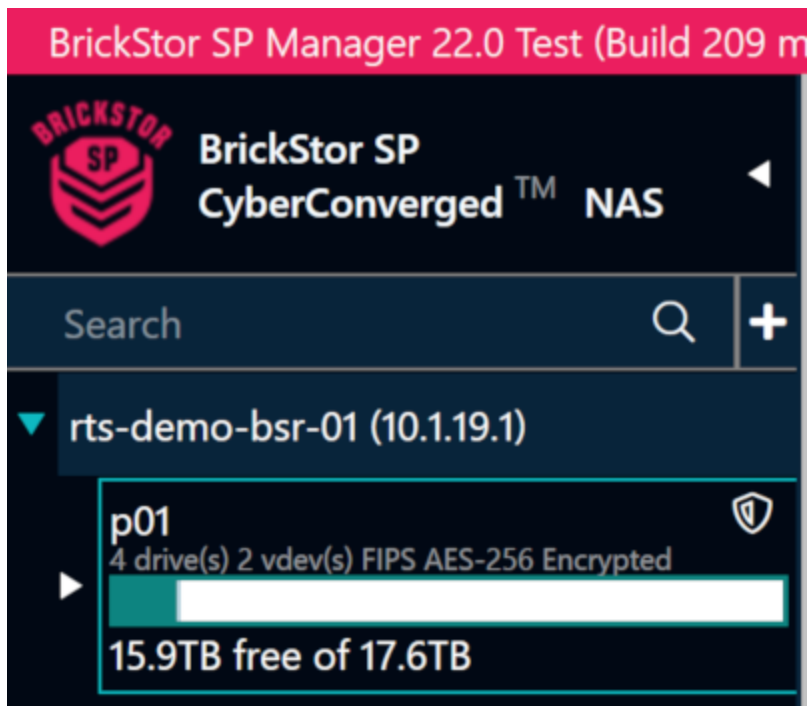
BrickStor SP Manager features several ways to modify pools that are currently on the system.

Expanding a Pool

There are multiple ways to expand a pool. The first is to select the pool in Rack View, select 'more' from the bottom bar, and then click any of the available expansion options.



The second option is to select the pool from the Connections pane on the left-hand side of BrickStor SP Manager and click either the Expand Data, Add Read Cache, Add Write Cache, or Add Spare button under the Pool heading, depending on what you would like to add to expand the pool (will only appear if the correct types of drives are available).



This will bring up the Expand Pool dialog box where you can choose to expand the pool by adding more vdevs, read and write caches, or spares. When the desired settings have been configured, click create to queue the change.

Expand Pool
Advanced

p01

Type

mirror ▼

Auto choose drives from alternating enclosures

Drive Type

107.4GB virtual VMWARE ▼

-

1

+

vdevs

-

2

+

drives per vdev

-


0

+

spare drives

Create
Cancel

All changes in the queue will be indicated in Rack View and must be committed using the changes tab on the right side of BrickStor SP Manager.



Growing a Pool


While expanding a pool primarily deals with adding additional disks to an existing pool, there's also the concept of growing the pool which is possible when the capacity of the underlying disk increases. This is typically possible when one of the following events occur:

- The pool is composed of mechanical or Solid State (SSD) drives and are replaced with a new ones of higher capacity.
- BrickStor SP is a VM and VM disk of the pool size is increased.
- The pool disk is an iSCSI or FibreChannel LUN and the size is increased on the underlying SAN solution.

Should the option to grow the pool become available following one of these events, do the following:

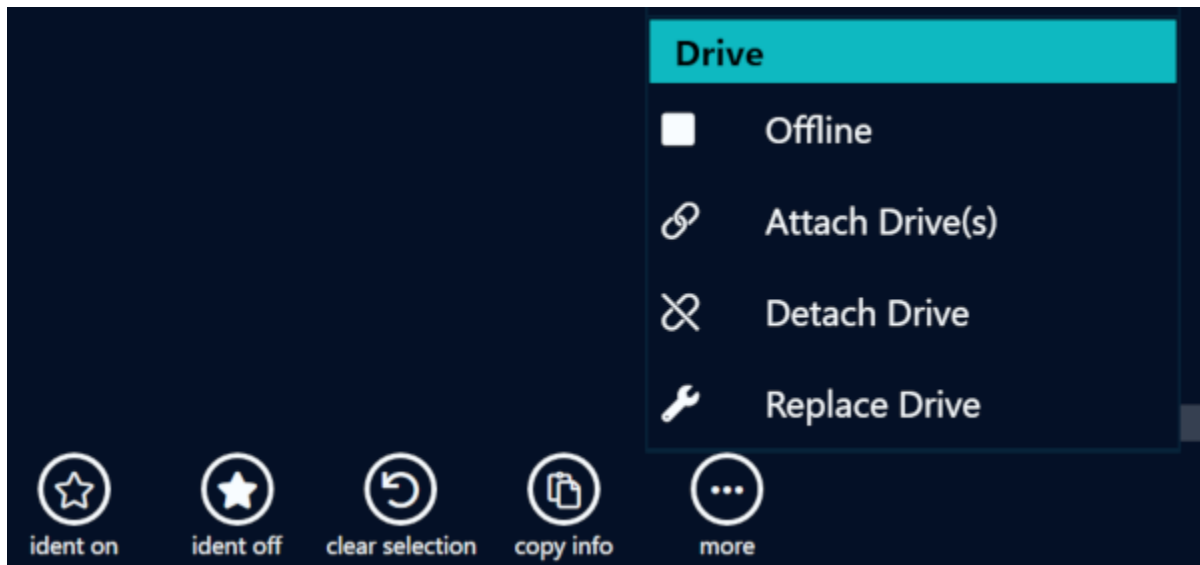
1. Using BrickStor SP Manager, select the desired pool.

2. Select the **General** tab.
3. Click the **Fix** button to grow the pool.




Replacing a Drive


If a drive becomes disabled or faulted it may be necessary to replace the drive with another available drive in the system. Select the drive you wish to replace in Rack View, click 'more,' and click 'Replace Drive'.




Or, if the drive is offline, you can navigate to the degraded pool in the Connections Pane on the left-hand side of the screen and click the Replace Drive button under the 'Notable Vdevs & Drives' heading.



Selecting an offline drive from Rack View will also bring up actions that can be performed on it.




This will bring up the Replace Drive dialog box where you can select the drive to use as the replacement then click the Replace button to queue the change.



The change will be indicated in Rack View and will not be committed until the Commit Changes button is clicked on the Changes tab.

Removing a Spare Drive

If a pool has a spare drive that no longer requires one, it can be removed to free up the drive by selecting the spare in the Rack View, selecting 'more,' and clicking the 'Remove Drive' button.




The change will be indicated in Rack View and will not be committed until you click the Commit Changes button in the Changes tab on the left-hand side.

Splitting a Mirrored Pool

A pool consisting of mirror vdevs can be split into two pools with no redundancy that contain the same data.

NOTE that this is only recommended in certain scenarios as the lack of redundancy increases the risk of data loss.

To split a mirrored pool, navigate to the pool from the Connections pane on the left-hand side and click the Split Mirrors into New Pool button under the Pool heading (you will need to click the arrow button to the right of the Pool heading to access this).



From the changes tab on the right-hand side you can change the name of the new pool that will result from the split and commit the changes with the Commit Changes button (by default the new pool created this way will be exported).

Attaching a Drive to a Pool

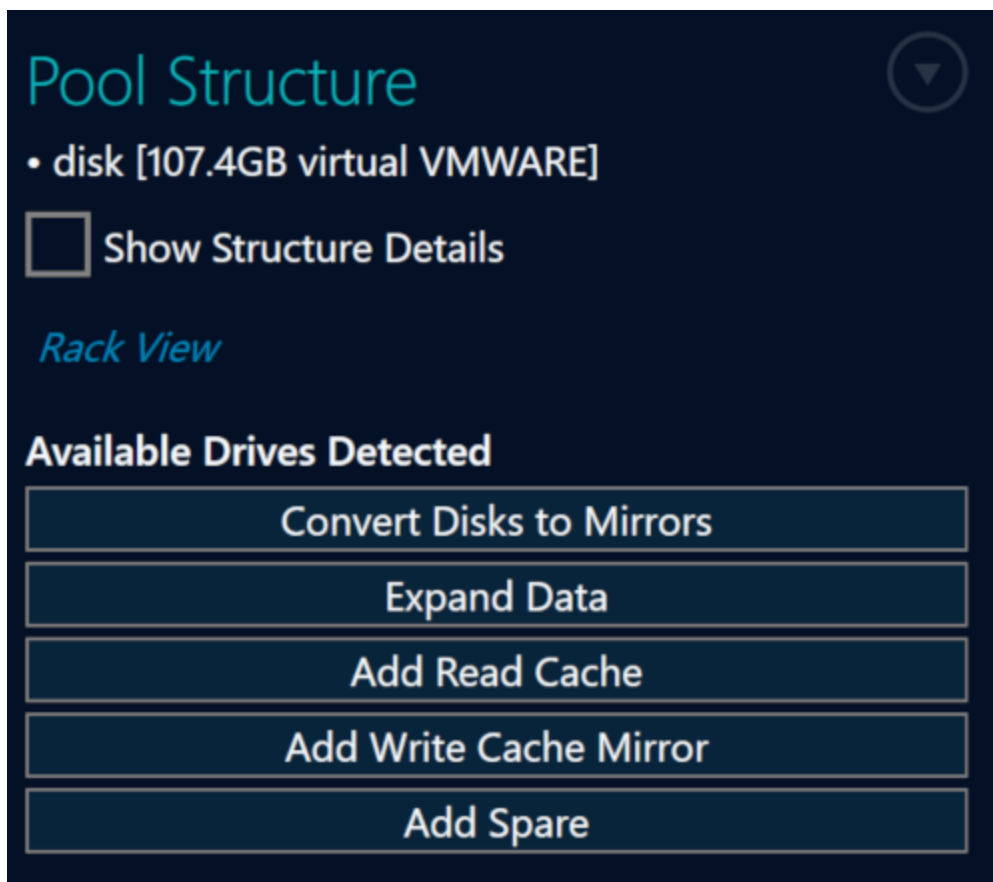
A pool with no redundancy can be converted to a mirrored pool, if there are enough available drives, in order to reduce the risk of data loss. To do this, select the pool in Rack View, select 'more', and click the 'Attach & Create Mirror' button.

The screenshot displays a context menu for a storage pool. The menu is organized into several sections:

- Destroy**: Contains a trash icon and the text "Destroy Pools".
- Pool**: Contains an upload icon and the text "Export", and a magnifying glass icon and the text "Start Scan".
- Available Drives Detected**: Contains a list of actions: "Convert Disks to Mirrors", "Expand Data", "Add Read Cache", "Add Write Cache Mirror", "Add Spare", and "Expand (Advanced)".
- Drive**: Contains a link icon and the text "Attach & Create Mirror", and a wrench icon and the text "Replace Drive".

At the bottom of the interface, there is a navigation bar with five icons and their corresponding labels: "ident on" (star in circle), "ident off" (star in circle), "clear selection" (refresh in circle), "copy info" (document in circle), and "more" (three dots in circle).

Or navigate to the pool from the Connections pane on the left-hand side and click the Convert Disks to Mirrors button under the Pool heading.



If done through Rack View, you will need to select the drive to attach yourself. When done through the pool's page it will select a drive for you automatically. The change will be indicated in Rack View and will not be committed until you click the Commit Changes button in the Changes tab on the right-hand side.

Trimming a Pool

If a pool is going to be retired or is no longer necessary and to be removed, it can be trimmed to the minimum operational set of drives. This will remove all redundancy and additional data protection and should only be done in specific scenarios. To trim a pool, navigate to the pool from the Connections pane on the left hand side and click the Trim Drives to Minimum Operational Set button under the Pool heading (you will need to click the arrow button to the right of the Pool heading to access this).



The steps it will take to trim the pool will be listed in the changes tab on the left-hand side and no changes will take effect until the Commit Changes button is clicked.

Changes 

Attach & Create Mirror undo
aaron-bsr1 (10.1.12.44)

 **test** [changes staged]
disk member
107.4GB virtual VMWARE

with

 **test** [changes staged]
disk member
107.4GB virtual VMWARE

Add spare undo
aaron-bsr1 (10.1.12.44)
spare [107.4GB virtual VMWARE]
to
test


Undo All

Commit Message

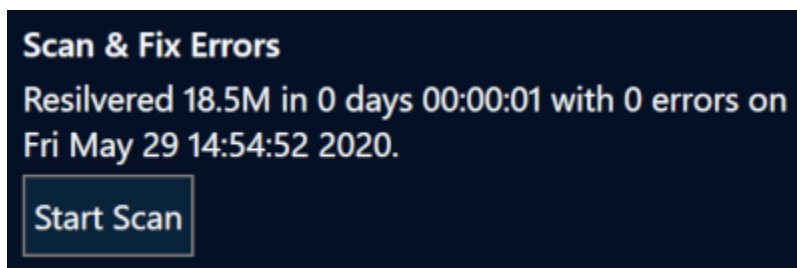
Commit 2 Change(s)

Scanning and Repairing a Pool

A pool can be checked for faults or problems and corrected using the scan pool feature. To scan a pool for potential faults, either select the pool in Rack View and click the more button at the bottom of the rack view and click Start Scan.




The button is also available on the Pool Tab.

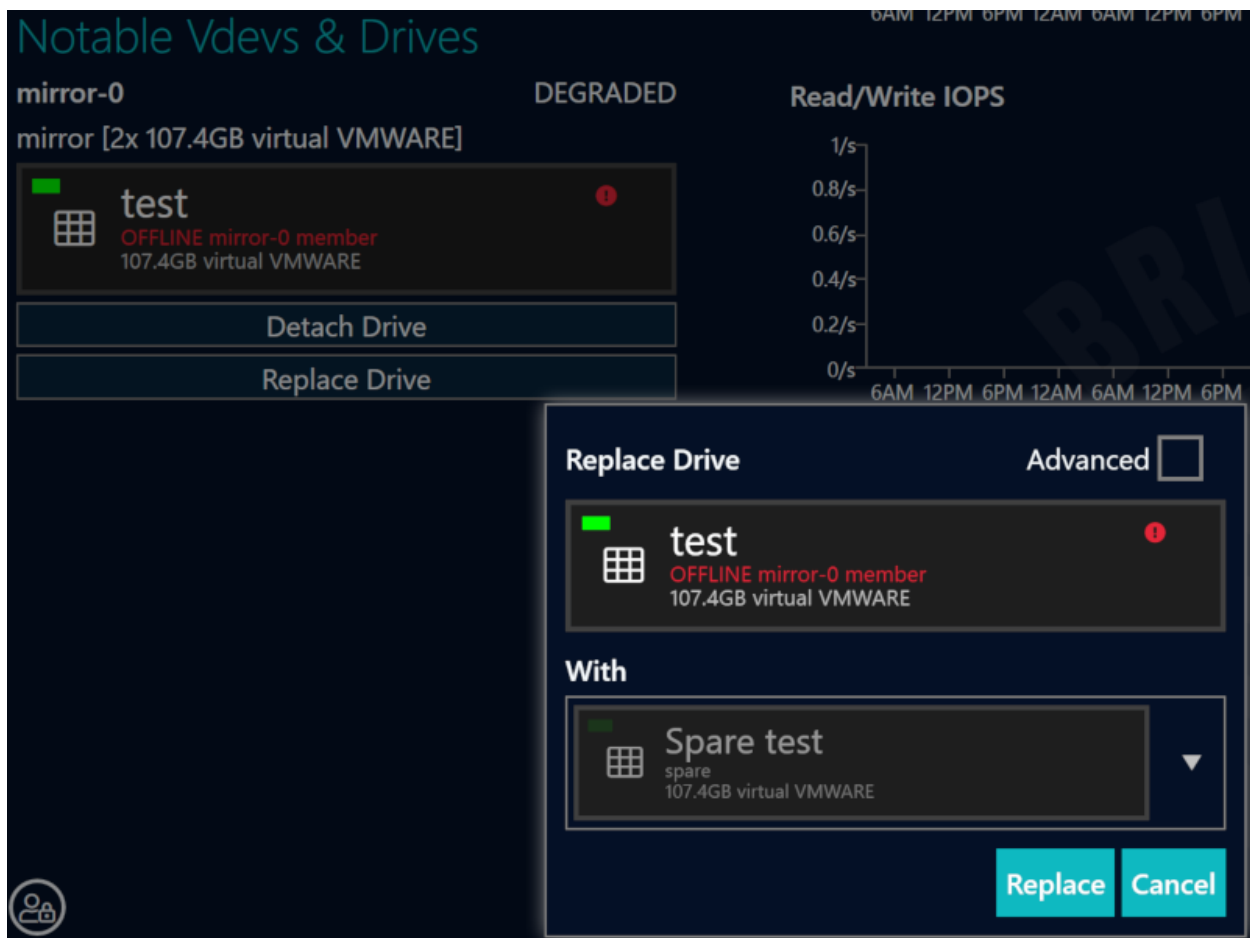


The scan will not be started until you click the Commit Changes button in the Changes tab on the left-hand side.

If the scan detects a faulty drive in the pool, it will mark the drive as degraded and replace it with a spare drive if one is available.



From the pool's screen on the Connections pane, the faulted drive will appear under Notable Vdevs & Drives. You can choose to promote the spare drive and detach the faulted drive from the pool, replace the faulted drive with another available drive on the system and return the spare to be a spare for the pool, or you can clear the errors on the drive if the problem has been corrected and return the spare. These options can also be found at the bottom of the screen in Rack View.



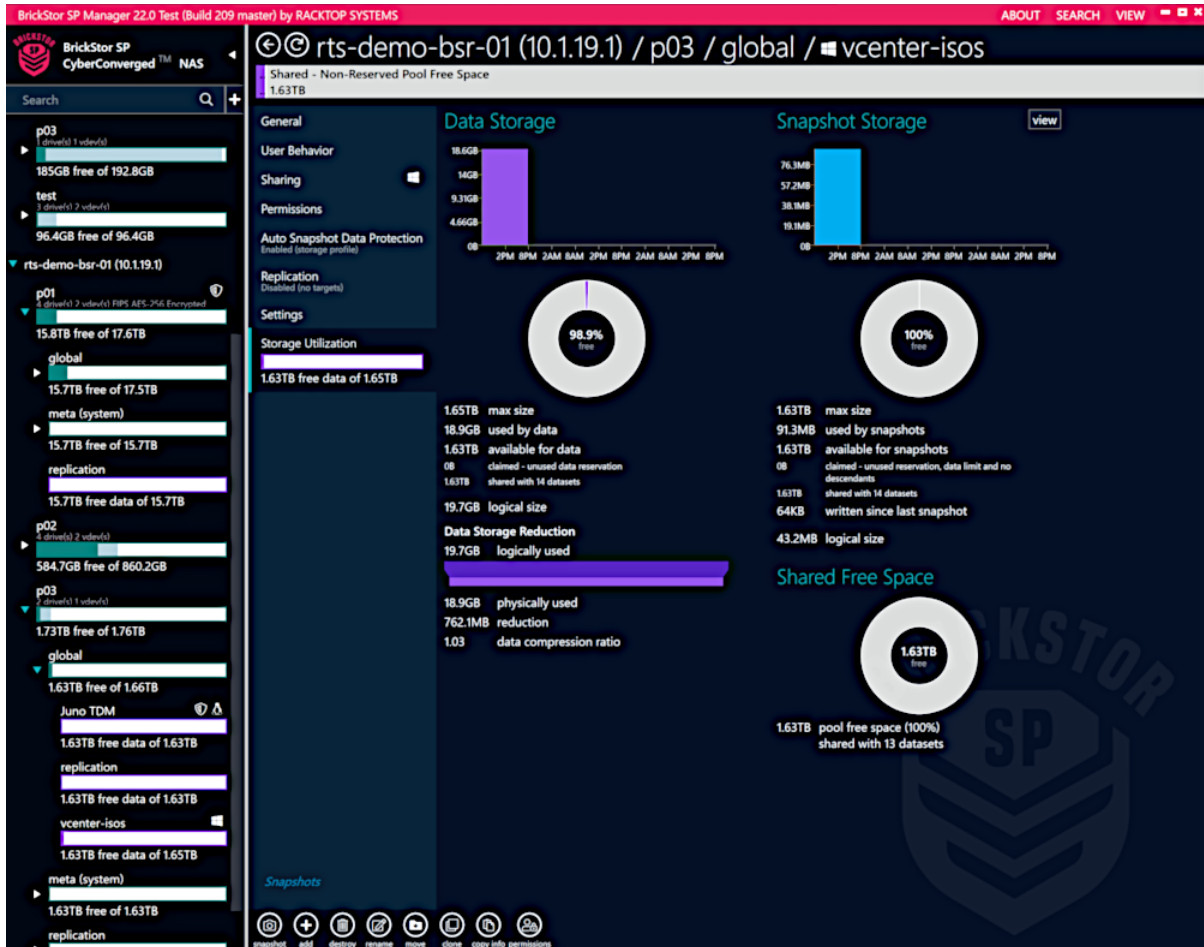
Each of these changes will require you to click the Commit Changes button in the Changes tab on the left-hand side to complete the action.

Pool Storage Utilization

Storage Utilization allows you to view information about the physical storage consumed by a pool.

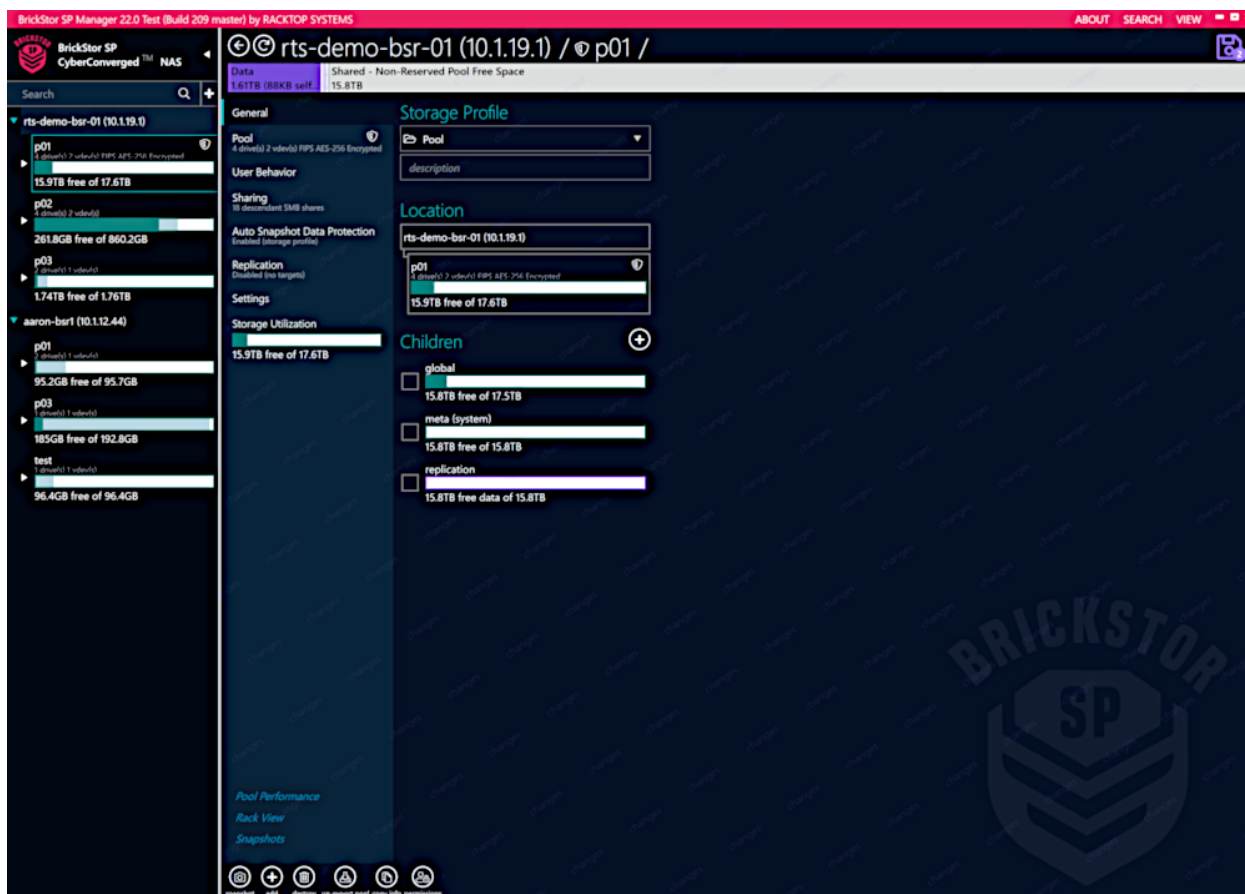
Viewing Pool Storage Utilization Statistics

1. In the Connections pane, select a pool.
2. In the Details pane, select Storage Utilization.



Pool Performance

Clicking on the 'Pool Performance' link leads to a page with charts and graphs about this pool's performance history.



Admins can zoom in on the graph to look at specific time periods.




Pool Sharing Information

The sharing tab shows the same information as the Sharing menu at the appliance level but scoped only to those shares on the selected pool.

Pool Settings


This tab contains settings that apply to the pool including a pool level reservation. The pool reservation by default is set to 10% of the pool capacity up to 100GB. This is in place as a safety measure to prevent the pool from becoming completely full and making it difficult to do the necessary operations to remove data. When the pool becomes full the admin can release some or all of the Pool reservation.

There is a hidden checkbox at the top of the page, 'show advanced,' that will provide more options.



Destroying Pools

To Destroy a pool, select the destroy icon while in the pool view. Once committed, this will destroy all descendant datasets and snapshots as well. You must double-click the pool(s) in the dialog to confirm.

**NOTE**

To ensure that all data is fully unrecoverable, there is also the option to Cryptographically Erase data on Self Encrypting Drives. This option is presented in the Changes pane during the commit. See [Cryptographically Erasing SEDs](#) for more details.

Datasets

Datasets are where you create and manage the file shares that end users use to complete their everyday work. After you have created one or more pools, you can create datasets within those pools.

Shares

Sharing from the dataset level is where the admins configure the share protocol and in the case of SMB and AFP the share name for the dataset.

Share Types

You can configure the following share types for your BrickStor storage.

- SMB
- AFP
- NFS

SMB

For SMB shares you have the option to enable the dataset to be shared out as a top-level SMB Share. If you enable Access Based Enumeration (ABE) the system hides the share from anyone browsing via SMB who doesn't have read access to that share. Host Base Access control further restricts access by source IP.

SMB Share

Connect Using

\\rts-demo-bsr-01\vcenter-isos

On

Hide from users that don't have permission (ABE)

Host based access control
Example: @1.2.3.*; @1.2.3.4/24; *.foo.com

Read-only

▼

Read/Write

▼

Deny

AFP

WARNING

AFP is deprecated in Release 22 and may be removed in a future release. RackTop and Apple recommend migrating to SMB.

Apple File Protocol uses Posix permissions and like SMB can use a different share name than the dataset name. To enable it as a time machine target you must enable it through BrickStor SP Manager first.

AFP supports host base access control like SMB.

NOTE

AFP is not supported on HA Resource Groups.

NFS

BrickStor supports NFSv3 and NFSv4.0/4.1/4.2. NFS 4 and above supports ACLs while the NFS v3 standard only supports host based access control and POSIX permissions. NFS shares must be the same name as the dataset and share the path of the dataset starting with /storage and then the pool name.

NFS Share

Connect Using

rts-demo-bsr-01:/storage/p03/global/vcenter-i...

On

Control access by specifying IP and hostname criteria below.
Example: @1.2.3.*; @1.2.3.4/24; *.foo.com

Read-only

Read/Write

Full Control (Root)

@10.1.19.*

Deny

Security Mode

local

Hide descendant datasets

Data security labels

With NFS v4.2 clients BrickStor will support context security labels when the Data Security labels box is selected

Clicking on the NFS Read/Write Volume will take you to performance metrics related to NFS and the dataset.

Creating Datasets

When creating a dataset, take note of the following caveats: * You cannot enable or disable dataset encryption after you have created the dataset and committed the changes. * You cannot disable deduplication for any dataset that has had it enabled without moving the data to a new dataset and destroying the old dataset. * Most other operations are reversible; however the changes only apply to new blocks and files as data in the dataset is modified and created.

To create a dataset, complete the following steps:

1. In the Connections pane, select either a pool or global container.
2. In the Details pane, click the add icon next to the Children label.

TIP You can also click the add icon in the lower portion of the Details pane.

The Create Dataset dialog box appears.

3. In the Create Dataset dialog box, type a name for the dataset.
4. Under **Type - Storage Profile**, choose one of the following options, based on your proposed workload:
 - If you are setting up a File System:
 - General File System
 - Rendering
 - Streaming Media File System
 - Apple Time Machine Storage
 - Archive File System
 - E-Discovery File System
 - Temp File System
 - If you're setting up Server Storage:
 - MongoDB Volume
 - MS Exchange Volume
 - Oracle Volume
 - If you are setting up Virtualization Storage:
 - Hyper-V Virtual Machines
 - Hyper-V Virtual Machines Volume
 - VMware VDI
 - VMware Virtual Machines
 - VMware Virtual Machines Volume
 - Xen Virtual Machines

- If you are setting up a Volume:
 - General Volume
 - Archive Volume
 - Temp Volume
 - If you are setting up a custom file system or volume:
 - Custom File System
 - Custom Volume
5. Select whether to enable **Dataset Encryption** on this dataset.

NOTE You must enable encryption during dataset creation.

6. Optionally, enter a **Data Quota**.
7. Accept the default **Data Reservation** or enter a new value.
8. Select your desired share type, either:
 - NFS
 - SMB
 - AFP
9. Click **Create**.
10. In the Changes pane, click **Commit Changes**.

Working with Datasets

After you create a dataset, BrickStor SP Manager allows you to modify most settings displayed in the initial create dataset dialog as well as additional settings.

Dataset Permissions

After you create a dataset, you can configure access control permissions for that dataset. When joined to Active Directory or LDAP you can use AD user names and groups. You can recursively apply permissions to a dataset and its descendants and reset ownership by selecting the appropriate check boxes.

Configuring Dataset Permissions

To configure dataset permissions, complete the following steps:

1. Select your dataset in the Connections pane
2. Select the Permissions tab in the Detail pane
3. Click the Add Permission button

Adding permissions to a dataset

BrickStor SP Manager 22.0 Test (Build 11 release/R22) by RACKTOP SYSTEMS ABOUT SEARCH VIEW

RackTop CyberConverged™ NAS

bsr-0fa3c2d9 10.1.12.102 Thu 4/23 by root

bsr-d6a77c45 (10.1.29.101)

poolparty 4 drive(s) 2 vdev(s) 17.8GB free of 18.4GB

global 16GB free of 16.5GB

Data2 16GB free data of 16.5GB

Data3 16GB free data of 16GB

Data4 16GB free data of 16GB

meta (system) 16GB free of 16GB

replication 16GB free data of 16GB

bsrqa02 10.1.29.109 Thu 7/9 12:59 PM by root

bsr-d6a77c... / poolparty / global / Data4

Shared - Non-Reserved Pool Free Space 16GB

General

User Behavior

Sharing

Permissions

Auto Snapshot Data Protection Enabled (storage profile)

Replication Disabled (no targets)

Settings

Storage Utilization 16GB free data of 16GB

TDM Disabled

File System Permissions

Full Control

Owner

List Folder Contents

Everyone

Read/Write

Domain Users@racktoplabs.com

Add Permission Copy From Reset

Add Permission

user/group name

Recently Used


- Domain Users@racktoplabs.com
- Everyone
- Owner
- Owner Primary Group
- root

Add Cancel

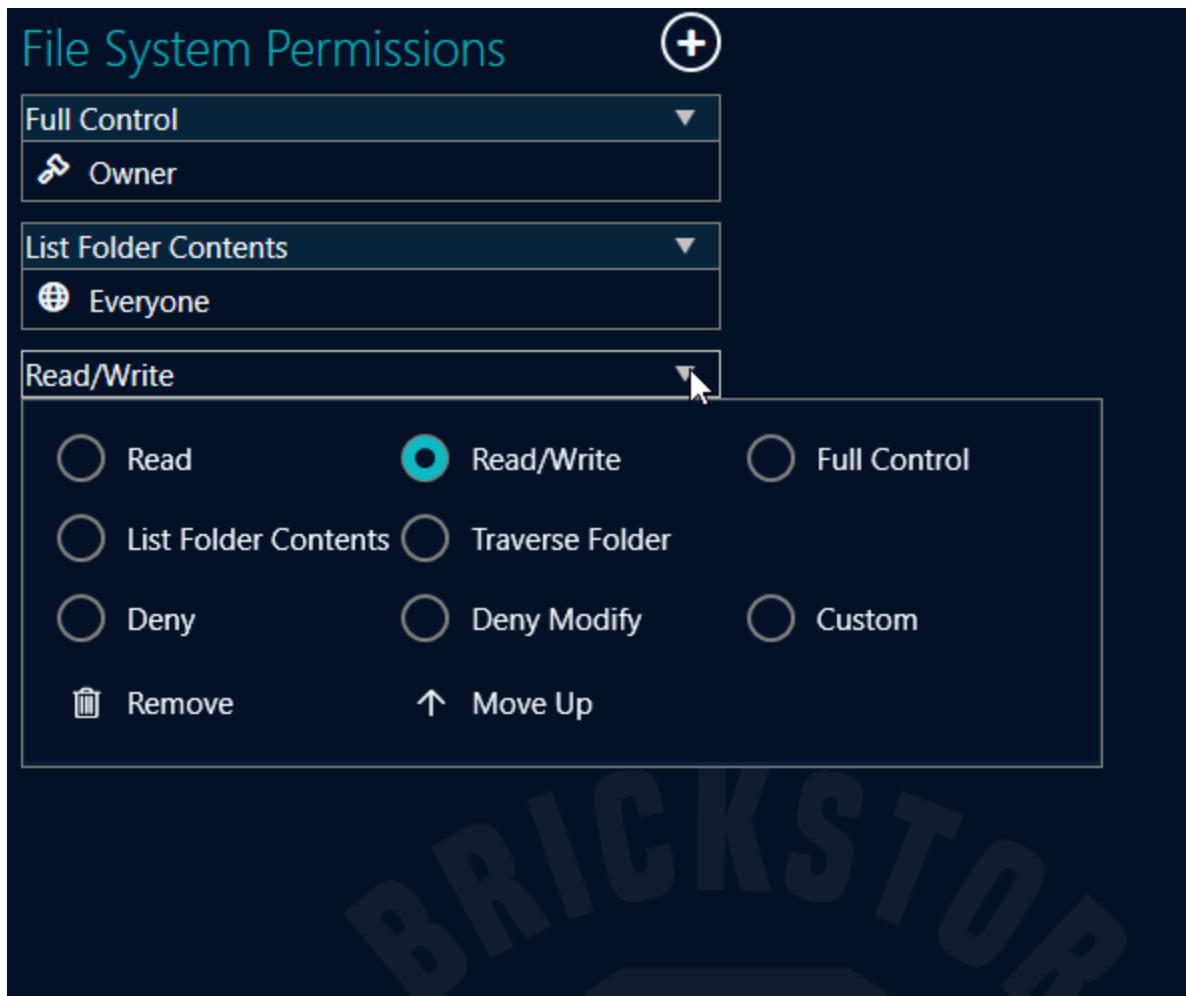
snapshot add destroy rename move clone copy info permissions

Using the Add Permission dialog, you can select previously used users or groups, or search for a user or group.

Add permissions search results




In the drop-down above the user or group, you can modify the type of permission. The default is Read/Write.



Additional options include recursively applying permissions or setting the new user or group as the owner. Once those choices are made, click the Commit button in the Changes pane to apply.

Choose permissions options and commit or undo



NOTE

If the Recursively Apply box is not checked, permissions will only apply to newly created files and folders. Files created in existing folders will not be updated.

WARNING

When Recursively Apply is checked, all files and sub-datasets will have permissions overwritten. On datasets with a large number of files, this operation could take some time as each file and folder is updated.

Copy Permissions from Another Dataset

Admins can copy the permissions of another dataset to the selected data set with the Copy From button. This feature will allow you to copy the permissions of any dataset on any appliance you are currently logged into.

Quotas and Reservations


After creating a dataset, you can configure quotas and reservations. You can quota only the data or you can quota the data with snapshots and descendants. You can also set reservations on the dataset for both instead of thinly provisioning the dataset. You can type a number and scale such as MB, GB, TB or you can use the slider above the text box to set the quota or reservation.

Dataset Storage Utilization

Storage Utilization allows you to view information about the physical storage consumed by a dataset.

Viewing Dataset Storage Utilization Statistics

1. In the Connections pane, select a dataset.
2. In the Details pane, select Storage Utilization.



iSCSI

BrickStor allows you to configure iSCSI targets. iSCSI targets are used by iSCSI initiators to establish a network connection. The target includes LUNs, which are collections of disk blocks accessible via the iSCSI protocol over the network. A target can offer one or more LUNs to the iSCSI clients that initiate a connection with the iSCSI server.

The system creates iSCSI volumes under the **Global/VBD** dataset.

In an HA cluster, iSCSI volumes fail over gracefully as part of the pool and resource group to which it was assigned. HA only supports iSCSI for boot devices.

Configuring iSCSI Volumes and Sharing as a Target

To configure a volume and share as an iSCSI target, complete the following steps:

1. SSH into the BrickStorOS as root.
2. At the BrickStor CLI, enter the following command to enable the target service.

```
svcadm enable -r svc:/network/iscsi/target:default
```

3. Enter the following command to create the default target.

```
# itadm create-target
```

4. Now, check the status of your targets to make sure they were properly configured, by running the following command:

```
# itadm list-target |v
```

```
TARGET NAME STATE SESSIONS iqn.2010-03.com.racktopsystems:02:c434c8d7-5643-6364-af5d-cb0bae33d531 online 0 alias: - auth: none (defaults) targetchapuser: - targetchapsecret: unset tpg-tags: default
```

5. Open BrickStor SP Manager and log into the BrickStor appliance to complete the iSCSI configuration.
6. In the Connections pane, select a Pool and then select the **General** tab in the Details pane.
7. In the lower portion of the screen, click the **Add** icon.

8. In the Create Dataset dialog box, type a name for the dataset.
9. Under Type-Storage Profile, select one of the following options:
 - General Volume
 - Archive Volume
 - Temp Volume
10. Select a Size, either using the slider or by entering a number.
11. Select a **Block Size**.

The dataset block size must match the block on the initiator's OS when you format the volume.

12. Check **Thin Provision** if you want to allocate disk storage space in a flexible manner, based on the minimum space required at any given time.
13. Under **Enter initiator(s) to share with**, type the name of the initiator.

TIP You can add multiple initiators in this field.

The initiator must be entered in one of the following formats:

- iqn: iqn.yyyy-mm.reverse-domain-name:unique-name
 - wwn: wwn.01234567ABCDEF
 - eui: eui.01234567ABCDEF
14. Under LUN, leave the field blank if you want the system to auto select the LUN that it will allocate.

To manually select a LUN, enter a value.

15. Click **Create**.
16. In the Changes pane, click **Commit Changes**.

Managing iSCSI Volumes

After you create an iSCSI volume, you can manage the volume on the Pool level Sharing tab in BrickStor SP Manager.

To manage iSCSI volumes, complete the following steps:

1. In BrickStor SP Manager, select the Pool level in the Connections pane.
2. In the Details pane, select an iSCSI volume under Descendent iSCSI volumes.
3. On the iSCSI page, you can complete any of the following actions:

To...	Do this...
Enable or disable an iSCSI volume	Click the toggle switch to either Online or Offline .
Delete an initiator	Click the adjacent trash icon.
Add an initiator	Click Add Initiator .
Remove initiators	Click Remove All .
Restore initiators	Click Restore All .

Encryption and Key Management

Managing Encryption

This tab shows the status and options relating to Self-Encrypting Drives (SEDs) and the Key Manager used for individual dataset encryption. Note that SED management requires a valid TCG license. For the Drives you can view which drives are SED capable. The boot pool is typically not SED capable or enabled.

SED Pool Status Meanings

- Not encrypted
- FIPS AES-256 encrypted
- FIPS AES-256 encrypted (data only) – Cache drives aren't SED
- FIPS AES-256 encrypted (partial) – Some data drives aren't SED
- FIPS AES-256 encrypted (partial enrolled) – Some drives have not been enrolled but are SED Capable

The screenshot displays the 'Key Manager' interface with the following sections:

- General:** Includes Sharing (24 SMB shares, 5 NFS shares), Auto Snapshot Data Protection, and Replication.
- Encryption:** Shows 4 encrypted drives and 15 encrypted datasets.
- Key Manager:**
 - Racktop Key Manager (Online since Tue 6/2 2:35 PM)
 - Buttons: Export All Encryption Keys, Import Encryption Keys, Resync Encryption Keys with Peers
 - Auto send key material to: 10.1.19.2 (Online since 9:09 AM, all keys | best effort backup)
- Encryption Services:** Shows Encryption Services as active.
- Drive Encryption (SED):**
 - 4 Enrolled - FIPS AES-256 Encrypted
 - 4 Unlocked - Ready to Auto-Lock
 - 7 Not Supported
 - Buttons: Verify Keys, Rekey, Export SED Keys, Unenroll, Config (Advanced)
- Encrypted Pools:** Shows pool p01 with 4 drive(s) and 2 vdev(s) FIPS AES-256 Encrypted.
- Dataset Encryption:**
 - 15 AES-256 Encrypted
 - 15 Unlocked - Accessible
 - 14 Not Encrypted
 - Buttons: Dataset Encryption Report, Share Encryption Report
 - Encrypted Datasets list: p01/global/frank/HIPAA Data, p01/global/frank/Lab Test, p01/global/jonathanGF, p01/global/jonathanlab, p01/global/Josh, p01/global/KirbyDemoProject1, p01/global/sample discovery, p01/global/Traudt, p02/global/01 CEO, p02/global/02 Development, p02/global/03 Finance, p02/global/04 Sales, p02/global/finance, p02/global/finance/testsets, p03/global/Juno TDM

Drive Encryption Related Buttons

Verify Keys – Checks that the node has access to all the appropriate data drive unlock keys through

the configured key manager.

Rekey – Changes the data drive unlock key for the data drives by requesting a new key from the key manager and applying it to the SED drive.

Export SED Keys – Exports SED keys to a password protected file that will be saved to the machine running BrickStor SP Manager. This feature must be enabled in the secured service configuration.

Unenroll – Unenroll takes the drive out of the FIPS compliant configuration, sets the drive not to auto lock when power is removed and sets the data drive lock key back to a known default. This feature must be enabled in the secured service configuration. This can be used if you want to transfer the disk to another system without having to share the key. However, the drive will not be protected in transit. It is also a safe way to change from one key manager to another and not have to worry about managing keys through the transition.

Config Advanced – This is only for modifying how often the secured service is performing low level functions.

Key Manager Buttons

Export All Encryption Keys – Exports SED and dataset keys to a password protected file that will be saved to the machine running the BrickStor SP Manager interface.

Import Encryption Keys – Imports keys from a password protected file created by BrickStor SP Manager.

Encryption Best Practices

For Users with the Local Key Manager

1. Regularly export the keys from the local key manager and save them in a safe controlled location off the BrickStor. In an HA cluster export and import the keys from both nodes to the other node and then export the keys from one node for backup. This should be done any time new encrypted datasets are created.
2. Import dataset keys to remote systems that are replication targets for fast recovery
3. Do not enable automatic key rotation
4. Enable key import and key export
5. Do not enable crypto-erase unless this is something you will need to do as part of regular operations
6. Do not enable unenroll drives so that nobody except an admin who modifies the config first can allow that operation
7. Periodically review the drive status report and the dataset encryption report
8. Manually perform a rekey based on organizational polices for encryption key rotation
9. Test recovery of files on the replication target to verify access to data during a non-critical time

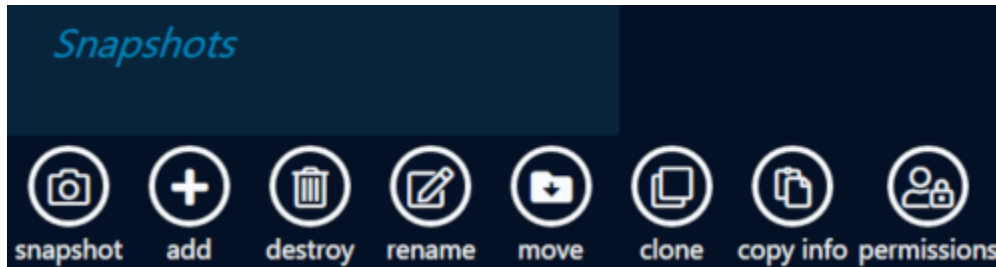
For Users with an External Key Manager

1. Verify your external key manager has appropriate backups and COOP plans.
2. Enable automatic key rotation
3. Determine if you want to enable key export based on your security posture and if you need them for COOP planning
4. Do not enable crypto-erase unless this is something you will need to do as part of regular operations
5. Verify replication targets can access appropriate dataset encryption keys on the key manager or export them and import them to the replication targets key manager.
6. Do not enable unenroll drives so that nobody except an admin who modifies the config first can allow that operation
7. Periodically review the drive status report and the dataset encryption report
8. Test recovery of files on the replication target to verify access to data during a non-critical time

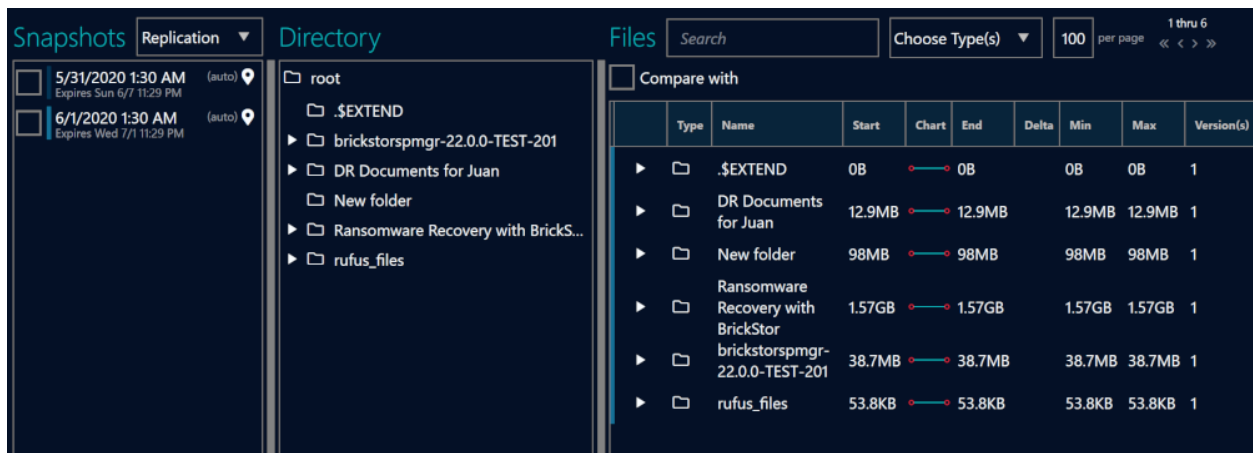
Snapshots

Snapshot Indexing

After snapshots have been created of a dataset, they can be accessed and viewed. After selecting a dataset from the Connections pane, click on the 'Snapshots' button near the bottom.




This brings up the snapshots screen for the selected dataset. At the top, filters can be set to only view snapshots from a certain time range. On the left, all the snapshots that match the filter parameters will be listed. Each one can be toggled on and off, and all the files present in the selected snapshots will be displayed in the panel on the right. Each file has a chart associated with it that shows its size over the course of the selected snapshots.




Restoring a File from a Snapshot

From the snapshots page, any item in any snapshot can be restored. To do this, click on the dropdown arrow on an item in the snapshot, and select 'Restore.'

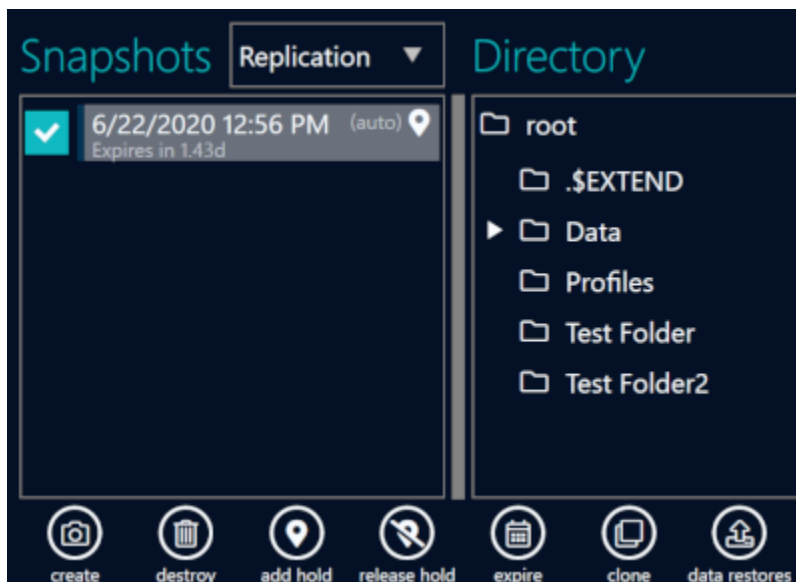


In the dialog box that shows up, choose whether the restored file should overwrite any existing file, rename the existing, or rename the restored file. Select 'Restore' to complete the action.



Snapshot Holds

It is sometimes necessary to hold snapshots past the normal expiration period. They can be assigned a tag that will be used to report on and enable an admin to remove all holds across all datasets on the appliance with that hold tag. You can also set an expiration on the hold tag itself. No snapshot will be removed from the dataset if there is a hold tag applied.

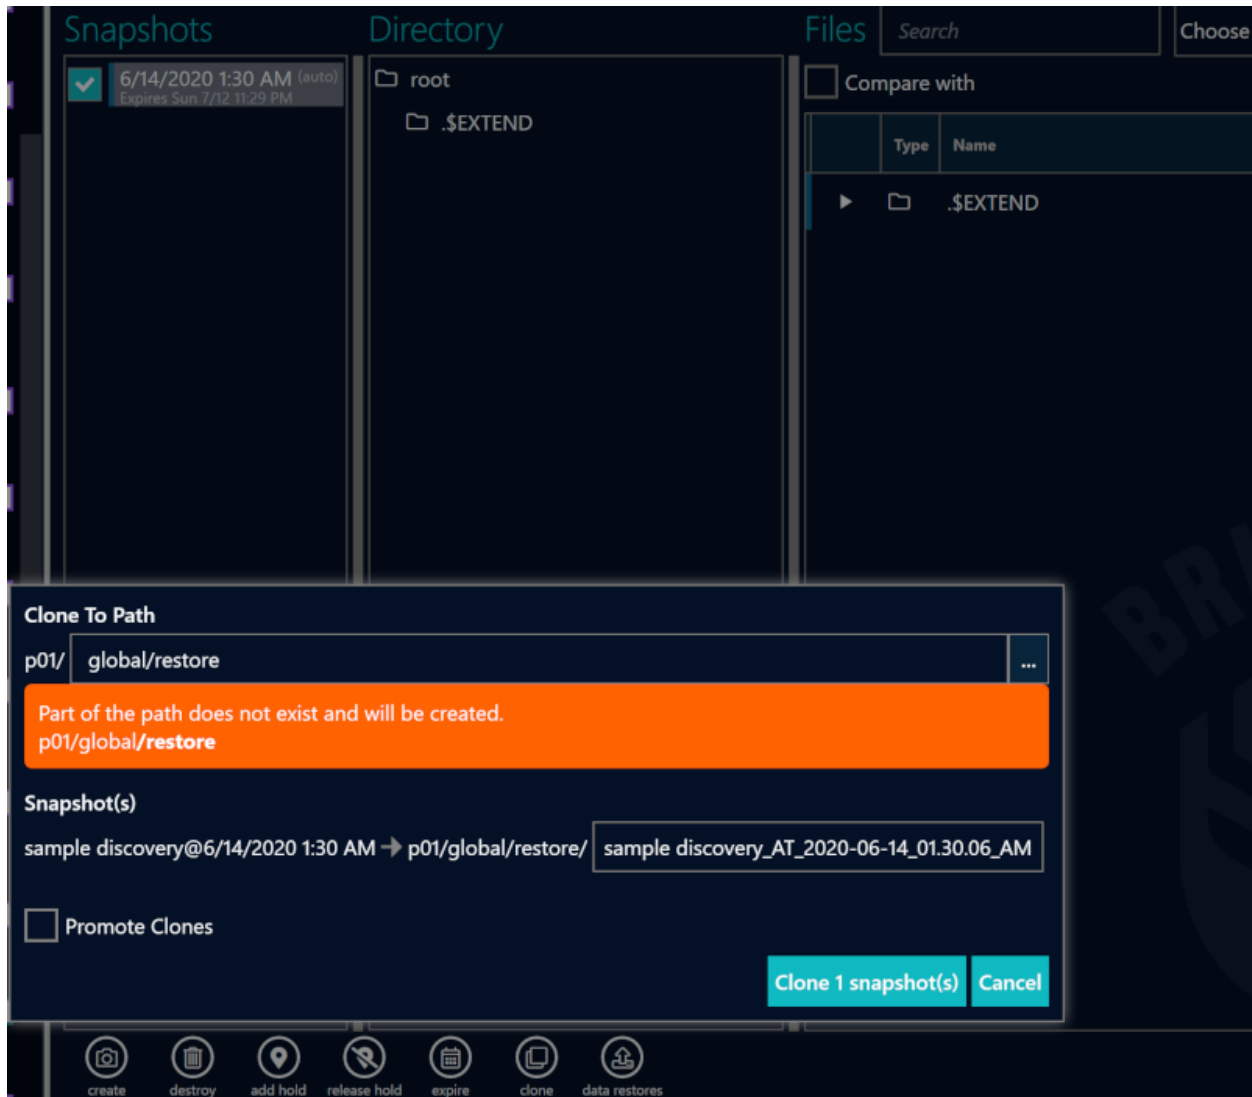


To release a hold tag you can just click release hold on the appropriate data sets.

If you delete a dataset you will delete the snapshots with it. If there are snapshots with a hold tag in the dataset pending destruction it will ask you to remove and release the holds before it can proceed destroying the dataset.

Clones

BrickStor SP allows you to select a snapshot to clone, which will create a writeable version of the snapshot without modifying the snapshot. Only changes to the clone will take additional capacity on disk. You can choose the path to create the clone. It must be on the same pool as the snapshot. Clones are the way to retrieve a file or files out of the snapshot on a replica because they are not mounted.



Be careful when promoting a clone. You should only promote a clone when you want all the snapshots prior to the snapshot to be linked to the clone and not the original active dataset. This operation is not reversible. It may also break replication if done improperly and you lose the common snapshot between the original and the replica.

Clones are a rapid way to create an entire dataset based on a point in time. This is a common method used to recover from a ransomware attack. They can also be used to create a version of a dataset to test an upgrade or run destructive tests and analysis against data without affecting the golden copy of data.

Replication

Data Protection includes integrated WAN optimized replication. BrickStor supports block and file level replication. Only the changed data is transmitted to shorten replication windows and reduce bandwidth usage. BrickStor replication supports bandwidth throttling. BrickStor Replication supports pause and resume as well as resume from bookmarks when interrupted by network outages and disruption. BrickStor supports block level replication to other BrickStor devices as well as file replication to any NAS or qualified S3 object storage. RackTop's data replication and backup capabilities enable customers to take advantage of a hybrid cloud strategy and use the cloud provider of their choice.

Replication Best Practices

1. When setting up replication, especially for larger data sets where data is being written, snapshots should be set to run more frequently than you may run them during normal operation. Each snapshot becomes a replication job, and since more frequent snapshots will be smaller, there is less likely to be a failure to replicate due to network errors or latency. Any replication retransmits are also more likely to be successful.
2. In cases where an encrypted data set is being replicated, keys should be exported from the local BrickStor and imported on the remote BrickStor so that the data can be recovered there.
3. Use the advanced configuration parameters to optimize your replication:
 - Priorities can be set to determine which data sets will replicate first
 - Bandwidth throttling can be configured to optimize how much bandwidth is used and at what times of day, so that you can take advantage of low traffic periods and avoid high traffic periods.
 - Optimize snapshot retention periods on both ends
 - On the local system, make sure that snapshots are not aging out before they are replicated.
 - On the remote system, you may want longer retention periods, but this will also consume storage, so consider this balance.
4. Replication peers should be on an appropriate data network that will be available and not interfere with other network traffic.
5. Setup bsradm notify for snapshot reporting so that you can be sure your replications are successful.

Understanding Peers

BrickStor supports block replication between two or more pools within the same system or across systems. To set up replication between two systems you must establish a peer relationship with the target system from the origin system. Once the peer relationship is created you can set up replication between pools on a per data set basis.

Configuring a Peer Relationship

To configure a peer relationship, complete the following steps:

1. In the Connections pane, select the appliance level.

2. In the details pane, click the Data Protection tab.
3. Click on the Add Peer Button at the bottom left of the details pane.

TIP

If peers already exist on your system, you can click the Add Peer icon next to the Replication Peers label.

Add Peer

Username/password - bidirectional

Username/password - one way

Pairing key

Hostname or Address

Credentials

Name

Description

Overwrite existing

Add **Cancel**

4. In the Add Peer dialog box, enter an IP address or hostname for the desired peer.

NOTE

Replication 2.0 now supports replicating to an HA cluster through the resource group. This will allow replication to continue operating even after a fail over. You only need to peer once to the resource group; The BrickStor OS will coordinate sharing keys between the cluster nodes. If you are replicating to an HA cluster be sure to use the destination resource group's address (VNIC) in this step.

5. Enter the username and password for the desired peer.
6. Click Add Peer.


The added peer appears under the Replication Peers label. The new peer will remain greyed out until you have added a target to that peer. You must repeat this process in order to replicate in the reverse direction on the other host.

Understanding Peer Status

The following table describes peer status messages that you may encounter.


Table 3. Peer Status


This...



Means the peer is...

Healthy No Backlog

This...	Means the peer is...
	<p>Configured without replication targets enabled for Peer</p>

This...	Means the peer is...
	<p>Unreachable and has a Problem, such as the target pool is not imported and will show up as [unk] or the target pool is out of space.</p>


Data Protection Replication

Data will be replicated to the target pool under the Replication Container. Through the GUI the source Hostname and IP will be visible along with the original dataset name. However, this information is stored in file system metadata on the replication target so it will not match the exact path name if an admin is browsing the file system on the pool.

Data Replication Hierarchy on File System

<Pool Name> - global - - replication o <Serial Number of Source BrickStor> □ Data Set GUID of Source Data Set

Data Protection Policy Configurations



Configure the Data Protection Policy for a Storage Profile

Managing Replication Details

You can manage replication details for a peer from the Replication Details page, to include:

- Set replication window settings for bandwidth throttling and peak business hours
- View and configure replication targets
- Enable/Disable targets
- Set inheritance (whether to inherit replication parameters from the parent)
- View timing and transfer status
- Export a replication report
- Show the history of replication jobs by clicking the Open History button

Accessing the Replication Details page

Clicking on a Peer's IP address will navigate you to the replication details page.

Status	Dataset	Location	Common	Pending	Size
In-sync	p03/replication Pool Replication Container	replica localhost			
In-sync	p03/replication/CN000001 Tenant Replication Container	replica localhost			
In-sync	p03/replication/CN000001/434796703698098510 Tenant Replication Container	replica localhost			
In-sync	p03/replication/CN000001/SN: ZZ0000S1/p01 Tenant Replication Container	replica localhost			
In-sync	p03/replication/CN000001/SN: ZZ0000S1/p01/glo Replica	replica localhost	@5/6/2020 10:34 AM		6.12GB
In-sync	p03/replication/CN000001/SN: ZZ0000S1/p01/glo Replica	replica localhost	@5/6/2020 10:34 AM		6.12GB

Replication Transfer History

You can view the details of transfers. This list can be filtered and exported. Details include:

- Time
- Duration
- Source / Destination
- Size
- Speed
- Success Status

The screenshot shows the 'Replications : Dataset Replication' tab in the BrickStor SP Manager. The interface is divided into several sections:

- Scope:** habsr02 (10.1.12.118)
- Status:**
 - In progress
 - Completed
 - Failed
 - Canceled
- Destination:** any destination
- Related Reports:** Dataset Replicas

The main table displays the following data:

Timestamp	Duration	Src
5/28/2020 4:00:28 AM -04:00	126.1ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-159276875122
5/28/2020 4:00:28 AM -04:00	152.7ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161371035848
5/28/2020 4:00:32 AM -04:00	16ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161632607295
5/28/2020 4:00:33 AM -04:00	139.4ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-38322832234
5/28/2020 4:00:35 AM -04:00	0.537s	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161371035848
5/29/2020 4:05:28 PM -04:00	124.8ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-159276875122
5/29/2020 4:05:28 PM -04:00	133.8ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161371035848
5/29/2020 4:05:33 PM -04:00	359.4ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161632607295
5/29/2020 4:05:34 PM -04:00	130.6ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-38322832234
5/29/2020 4:20:13 PM -04:00	194.4ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-159276875122
5/29/2020 4:20:18 PM -04:00	168.7ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161371035848
5/29/2020 4:20:18 PM -04:00	91.1ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-38322832234
5/29/2020 4:20:18 PM -04:00	117.8ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161632607295
5/29/2020 4:39:48 PM -04:00	145.2ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-159276875122
5/29/2020 4:39:48 PM -04:00	150.3ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161371035848
5/29/2020 4:39:53 PM -04:00	181.5ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-38322832234
5/29/2020 4:39:53 PM -04:00	147.2ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161632607295
5/29/2020 4:58:02 PM -04:00	179.1ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-161371035848
5/29/2020 4:58:02 PM -04:00	159.9ms	p02/replication/local/CN000001/8307030393173412241/I-8307030393173412241-159276875122

Auto Snapshot Data Protection

This tab allows the user to change the data protection policy from the default storage protection profile to a custom data protection policy specifically for that dataset. It is also the menu in which the admin can choose the replication target(s) for the dataset based on available peers. Replication targets can be inherited so if this is a dataset where global has already been defined with a replication target or in the case where it is a subdirectory with a defined replication target this will show up in the Auto Snapshot Replication portion of the tab.

The screenshot displays the BrickStor SP Manager interface. The top navigation bar shows the application name and version: "BrickStor SP Manager 22.0 Test (Build 209 master) by RACKTOP SYSTEMS". The breadcrumb path is "bstr-fa6e11e1 (10.1.19.2) / DRPOOL / global / DemoCluseISOimages". The main content area is titled "Auto Snapshot Creation" and shows the following settings:

- General:** 12.3GB written since last snapshot.
- Sharing:** Use profile protection policy.
- Permissions:** Use profile protection policy.
- Auto Snapshot Data Protection:** Enabled (storage profile).
- Replication:** Disabled (no targets).
- Settings:**
 - Frequency:** Every 4 hour(s)
 - Retention:** 1 day(s)
 - Daily consolidation:** 14 day(s)
 - Weekly:** no week(s)
 - Monthly:** no month(s)
 - Yearly:** no year(s)
- Storage Utilization:** 6.64TB free data of 6.65TB
- TDM:** Disabled
- Auto Replicated Snapshots:** Have same retention
- Recent Data Restoration:** Pending (0), Failed (0), Succeeded (0)

The sidebar on the left shows a tree view of storage targets:

- aaron-bsr1 (10.1.12.44)
- rts-demo-bsr-01 (10.1.19.1)
- bsr-fa6e11e1 (10.1.19.2)
 - DRPOOL (2 drives(s) 1 volume(s))
 - 6.74TB free of 7.04TB
 - global (6.64TB free of 6.65TB)
 - DemoCluseISOimages (6.64TB free data of 6.65TB)
 - meta (system) (6.64TB free of 6.64TB)
 - replication (6.64TB free of 6.93TB)
 - SSD/vmware (3 drives(s) 1 volume(s)) (860.5GB free of 860.5GB)
- habsr02 (10.1.12.118)

The bottom navigation bar includes icons for snapshot, add, destroy, restore, mount, close, copy data, and permissions.

When replication is enabled the screen will look like the below screen capture. By clicking on the [SNAPS] button under replication targets the GUI will take you to the snapshots page on the target node as long as you have network access from where the BrickStor SP Manager GUI is running.

You can remove targets and replace targets from this screen.

Clicking on the target will take you to a window that will show you all datasets on the appliance that are replicating to that target.

Auto Snapshot Creation log

12.3GB written since last snapshot

Use profile protection policy ⓘ

Automatic snapshot frequency and retention has been optimized for the selected storage profile. To edit these settings, choose custom.

On

-1 year
-6 month
now

Frequency		Retention			
Every 4 hour(s) ▼	-	1	day(s)	+	
Daily consolidation	-	14	day(s)	+	
Weekly	-	no	week(s)	+	
Monthly	-	no	month(s)	+	
Yearly	-	no	year(s)	+	

These settings only apply to new snapshots. Existing snapshots will expire based on the settings at the time of snapshot creation. Sub-daily snapshots will be skipped when no change occurs.

Auto Replicated Snapshots

Have same retention ▼

Recent Data Restoration log

Pending (0)

Failed (0)

Succeeded (0)

User Behavior Auditing and Analysis

User Behavior Auditing allows you to track how end users interact with data stored on your system. User Behavior logs the operations for each file made by applications and users, such as file creation, movement, deletions, and so on. BrickStor displays this information in real-time reports and graphs.

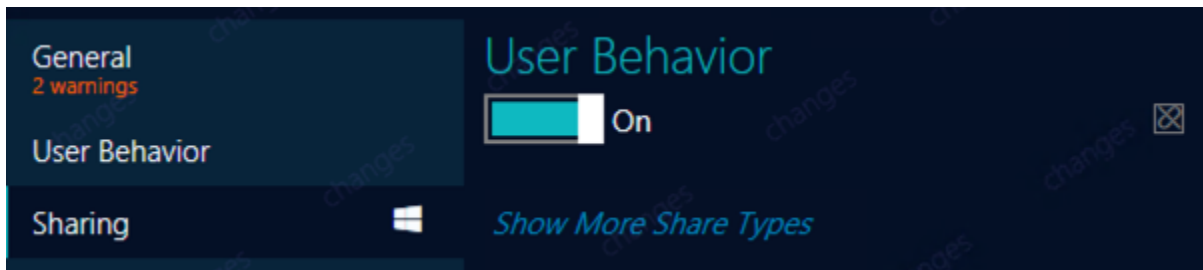
You can enable User Behavior at the pool level or the dataset level. BrickStor logs the behavior of users at the system level where it was configured and its descendants. For example, if you enable User Behavior at the Pool Level, it is also enabled for all datasets within that pool.

By default, the system stores user behavior data in the meta dataset of the pool.

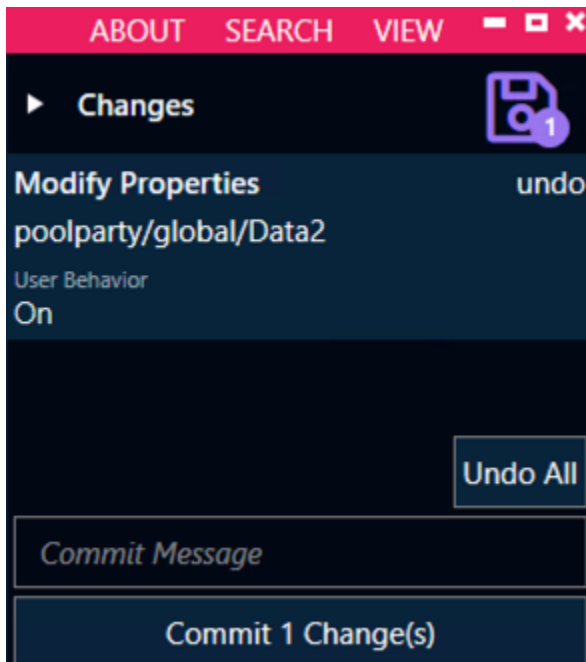
Enabling User Behavior

To enable User Behavior, complete the following steps:

1. In BrickStor SP Manager, select either a pool or dataset.
2. In the Details pane, select the **Settings** tab.
3. Under User Behavior, click the toggle button to **On**.



4. In the Changes pane, click **Commit Changes**.




User Behavior Audit

After you enable User Behavior, BrickStor displays an overview of all user actions initiated from that point. You can view the following information in the User Behavior Audit.

Accessing the User Behavior Audit

To view the User Behavior Audit, complete the following steps:

1. In the Connections pane, select either a pool or dataset.
2. In the Details pane, select the **User Behavior** tab.



Most of the content here can be clicked on and will lead to the *Activity* page.

Forwarding User Behavior

The user behavior activity can be forwarded to a SIEM or log centralization for off system processing and analysis. To configure UBA to forward to another host edit the configuration file in

```
/etc/racktop/ubcollectd/ubcollectd.conf [Syslog] Protocol = "udp" Server = "10.1.29.X:514"  
CertFile = "" Facility = "local0" Enabled = true
```

High Availability

To guarantee the highest level of data availability, the High Availability (HA) feature allows you to leverage an additional storage node to manage the underlying disk. Each storage node already comes built with all redundant hardware such as dual power supplies, multiple CPUs, two or more Host Bus Controllers (HBA), multiple network interfaces and so on. HA provides an additional layer of protection for other unforeseen system faults and zero-impact software upgrades.

BrickStor HA nodes operate as active/active so additional performance can be gained depending on the application.

High Availability Components

A BrickStor High Availability Cluster consists of four main components:

BrickStor Head Node – The Head Node is a hardware and software component responsible for managing underlying disk and presenting it as consumable data via SMB, NFS or iSCSI. BrickStor HA configuration consists of two Head Nodes communicating between each other with a shared configuration, system state and leverage a master election process.

Both nodes always have identical hardware configurations and operate on the same software version. Some versions are backwards compatible but only during the upgrade process. Please reference the release notes to find an upgrade path.

Heartbeat - Heartbeat is a method of Head Nodes communicating their health status. This is typically done over a dedicated network interface directly connecting both nodes. Additionally state is also communicated over the management interface "admin0". During complete loss of the node heartbeat the failover process will take place.

RMM/iLO - RMM is Intel's Remote Management Module and iLO is HPE's Integrated Lights-Out management facilities for out-of-band server access. Both are proprietary dedicated hardware components embedded on the motherboard to provide hardware management during the lights-out scenarios.

BrickStor HA relies on this interface during automated HA failover events to avoid split-brain situations. Split-brain is when heartbeat communications are compromised but both nodes are online and healthy.

Witness – The witness is an essential component for leveraging automated failover events. It is used to act as the third party in the quorum to break a tie. A witness is a software component that can either run Windows Server or Linux as virtual machine or a bare metal system. It installs as a lightweight service and communicates with both HA Head Nodes via the management interface.

The Witness does not take any part during manual failover initiated by a system administrator nor does it play any role in data presentation.

Shared Storage – Shared Storage refers to the underlying physical or virtual disk accessible by both Head Nodes.

Physical disk is presented with drive enclosures connected with redundant SAS connections to both nodes. It is highly advised to configure HA solutions with two or more enclosures and configure storage pool(s) with disks split across them. This ensures the solution can survive enclosure failure.

Virtual disk refers to block storage volumes presented to BrickStor HA Head Nodes by one or more third-party SAN solution(s). In those cases BrickStor HA is acting as an NFS/SMB protocol server consuming SAN volumes via iSCSI/FibreChannel links.

Storage Pool - A Storage Pool is an aggregation of physical or virtual devices describing physical characteristics of the storage system (capacity, performance and data redundancy). The pool is typically defined during system deployment and cannot be changed except to grow it by adding more devices. A given storage system can have one or more storage pools depending on the application. More on the storage pools can be found in [Storage Pools](#) section.

In an HA configuration only a single Head Node can serve a given pool. The second node would simply wait to take over (failover).

WARNING

Be advised, one should not attempt to import or export pools using the CLI. This will result in data corruption. Always use RackTop supplied utilities such as BrickStor SP Manager.

VNIC - A VNIC is a Virtual Network Interface which extends the functionality of a physical network port. VNICs are used by BrickStor HA to facilitate failover having data VNIC(s) float between the HA nodes.

WARNING

Use VNICs conservatively. Unusually large number of VNICs may affect failover times because each one must be reconstituted on failover.

Resource Group – A Resource Group is a logical grouping of Storage Pools and one or more VNIC(s). An HA Cluster can have one or more Resource Group and are typically created during solution deployment time.

Resource Groups can be modified, disabled, removed or moved between nodes. The following action can result in loss of data availability so use it with caution. Familiarize yourself with [Managing Resource Groups](#) before attempting to use them.

Resource Group Pool States

A pool within a Resource Group can be in one of five states when managing an HA cluster:

1. **Member of a Resource Group** – Pool is part of an HA Resource Group and is Enabled. The enabled pool is imported on the specified node and the second node is ready for failover.
2. **Disabled Member of Resource Group** – The pool is a member of a resource group but is administratively disabled. The disabled pool is exported from both nodes and data is not available. Once the Resource Group is enabled the pool will be imported on the specified HA node.
3. **Unmapped Pool** – Pool is a member of the HA Cluster but is not assigned into any current Resource Groups. This typically results when the pool is protected from being imported on more than one node at a time or brought over from a foreign HA configuration. In this state the pool is not imported on either nodes and can either be assigned into a Resource Group (new or existing) or destroyed.
4. **Removed from Cluster** – Pool is not a member of the HA configuration. In this state the pool is not imported on either node and can either be assigned into a Resource Group (new or existing) or destroyed.

5. **Missing** – The pool devices are not accessible by both HA nodes. This can result from the drives being physically removed from the enclosures, loss of connectivity with a drive enclosures or SAN, or the drives are SED (Secure Encrypted Drive) and are currently locked.

Standard Network Interfaces

At a minimum an HA configuration requires each node to have at least three physical network interfaces.

Management interface can also be referred to as "admin0". It is used for system management and HA communications.

Heartbeat interface directly connects each node and is used for exchanging HA communications between the nodes.

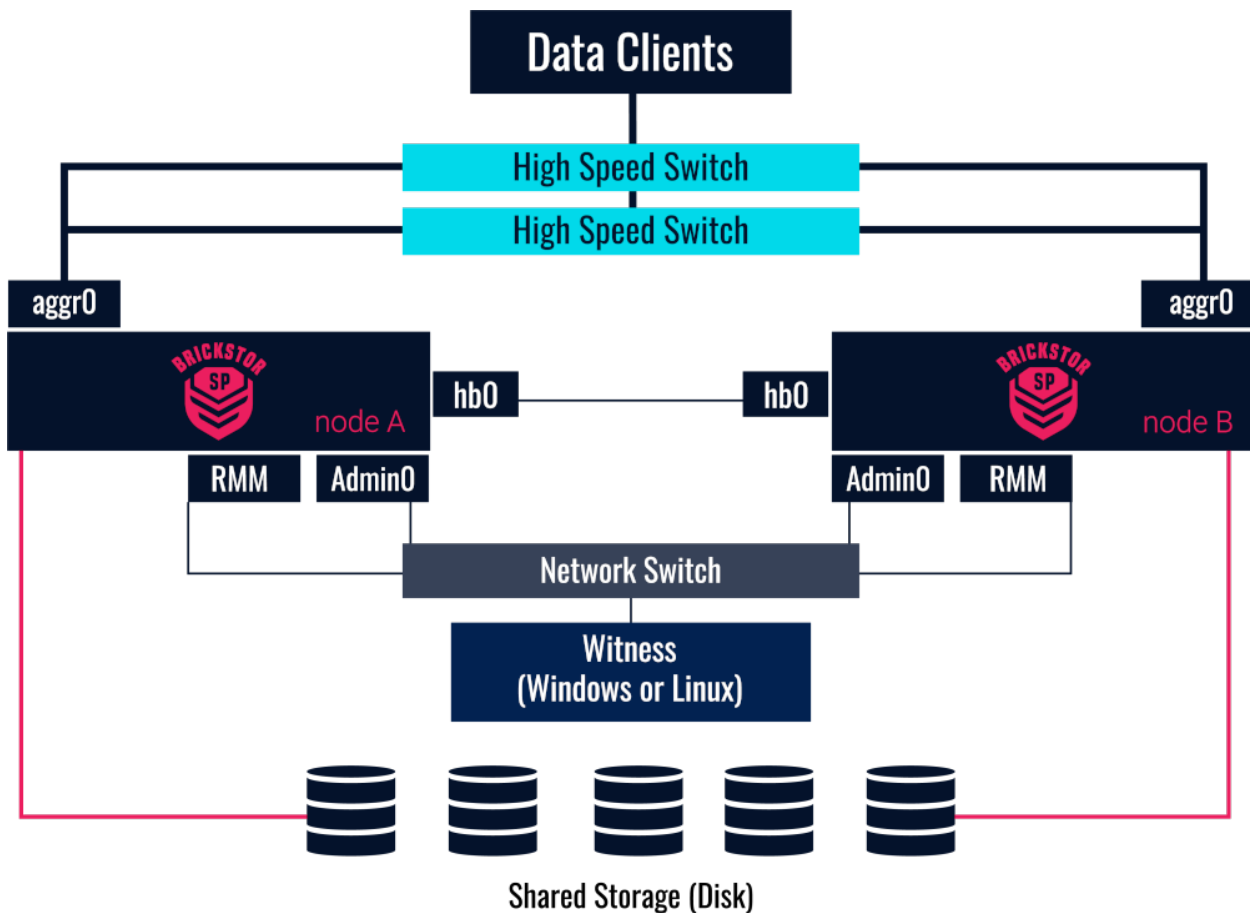
Data interface is client data access. This interface is typically composed of two or more physical interfaces aggregated together using LACP protocol (IEEE 802.3ad)

TIP It is highly advised to designate a secondary HA management interface over the data aggregate. This insures highly available network connectivity between the nodes and a witness server.

TIP The data aggregate interfaces should connect to two or more stacked high speed network switches.

WARNING The HA witness server and RMM/iLO interfaces must reside on the same subnet as the HA management interface.

HA Cluster Architecture



HA Scenarios

Loss of Management Network Connectivity

Loss of the HA node's management interface will prevent automated failover due to inability to communicate with the witness server. However, this will not directly impact the data availability. The system administrator would still be able to failover any Resource Groups using the second, healthy Head Node.

To overcome this edge case an additional management interface can be established over the data aggregate and designated for HA communications.

Loss of Data Network Connectivity

Loss of data network entirely is highly unlikely when it is configured as an aggregate of two or more physical ports. In this unlikely event or when Head Node is configured with only a single data interface the HA can be configured to failover on data interface loss.

Loss of RMM/iLO Connectivity

RMM and iLO communication for HA functionality is only used as a third means of power state verification. Loss of lights-out interface has no impact on system functionality given the management and heartbeat interfaces are healthy. If all means of communication are unavailable for a given node,

a failover event will take place.

Manual Failover

Manual failover is an action triggered by the system administrator to initiate a Resource Group(s) migration to the adjacent HA node. This action is typically performed as part of the solution maintenance or during upgrades.

Automatic Failover

Automatic failover takes place during HA node failure. In this event the surviving node will take over all the Resource Groups.

After automatic failover takes place it is advised to disable the failed node to prevent further action. For example: Resource Groups can be configured to have a preferred node. Disabling a failed node will prevent resources from failing over back and forth in the event this HA node is exhibiting an inconsistent behavior.

Once the issue is cleared up or the failed node is repaired it can be enabled again to return it back to service.

WARNING

The HA Witness server must be online, healthy and accessible by the surviving node in order for the automatic failover to trigger.

High Availability (HA) Best Practices

1. Use a dedicated witness for each HA cluster.
2. With HA witness being a VM be sure it is not running on the datastore using the same BrickStor HA shared storage.
3. Use two or more resource groups to get more performance out of your BrickStor making it active/active.
4. Use an LACP 802.1ad aggregate for the data network across two or more stacked network switches. This will boost network performance by load balancing traffic across multiple ports and improve availability.
5. Avoid manually failing over Resource Groups with pools in the degraded state Degraded pools can take longer to import during failover and this can result in a self induced outage. Resolve pool issues first and only then fail over the Resource Group.
6. Avoid using jumbo frames. Jumbo frames can boost performance for data transport. However, in NAS solutions it only fits in very specific environments and must be properly configured on all network devices. Improper use of jumbo frames can result in poor performance.
7. Avoid using DNS hostnames for HA configuration. This eliminates dependency on DNS services.

Configuring High Availability

Prerequisites

Before diving into cluster setup wizard the following prerequisites must be met in order to form a BrickStor HA cluster:

- All devices (2x HA Head Nodes and a witness server) must be properly connected and powered on.
- BrickStor SP Manager software installed and connected to both Head Nodes.
- Witness server:
 - `hiavd` service must be installed and running.
 - Must be able to ping both Head Nodes.
 - Must be able to connect via TCP port 4746 to each Head Node `telnet <node address> 4746`.
- Head Nodes:
 - Must be connected to disk enclosures with one or more disks present.
 - Data pool must be created and accessible by both nodes.
 - Heartbeat Ethernet port must be properly connected and configured.
 - Data aggregate must be created and working.

Once the following checks are completed you are ready to create the HA cluster using BrickStor SP Manager.

Setting up Witness Server

BrickStor HA Witness comes in the form of a single binary file shipped with each BrickStor system. It can be downloaded for either Windows Server or Linux by going to the web page of the BrickStor Appliance <https://<BrickStor Admin0 IP>:8443>.

The witness binary version must match the version of the HA Nodes. This process ensures one always has the correct binary for their deployment.

Installing Witness (Windows)

1. Retrieve a copy of the Windows `hiavd` executable by going to the web page of one of the BrickStor HA Head Nodes <https://<BrickStor Admin0 IP>:8443>.
2. Create a service home directory `c:\racktop`.
3. Extract the downloaded `hiavd.zip` into the `c:\racktop` directory.
4. Register as a Windows service
 - a. Open a command prompt or Powershell as an Administrator
 - b. Change directory to service home `cd c:\racktop`

- c. Install the service by typing `hiavd.exe /install`
- d. Configure the service to restart on failure by typing `sc failure "hiavd" actions=restart/60000/restart/60000/restart/60000/60000 reset=0`
- e. Start service by typing `sc start hiavd`

Configure Witness Firewall

The Witness service communicates via TCP port 4746 as well as ICMP protocol with the HA Head Nodes. The traffic must be allowed for both inbound and outbound communication on the witness server.

1. Open Windows Firewall configuration
 - a. Using Control Panel open Firewall `Control Panel\System and Security\Windows Defender Firewall`.
 - b. Select `Advanced Setting`. This will bring up a Windows Firewall Configuration window.
 - c. Select `Inbound Rules`.
2. Allow ICMP
 - a. From the rules list find and edit `File and Printer Sharing (Echo Request ICMPv4-In)`.
 - b. Using the `General` tab be sure `Action` is set to `Allow the connection`.
 - c. Using the `Scope` tab be sure `Remote IP Address` is set to `Any IP Address`.
 - d. Click `OK`.
3. Allow TCP port 4746 HA Head Nodes to communicate with the Witness service.
 - a. Using the `Action` menu select `New Rule...` to create a new inbound firewall rule.
 - b. In `Rule Type` select `Port` type
 - c. For `Protocol and Port` use `TCP` and for `Specific local ports` enter 4746.
 - d. For `Action` select `Allow the connection`.
 - e. For `Profile` select all available profiles or choose ones that apply to your environment.
 - f. For `Name` enter a meaningful name such as `RackTop BrickStor HA Witness TCP 4746`.
 - g. Click `Finish`.

TIP

When Antivirus software is installed on the witness server be sure to exclude `hiavd` service home directory `c:\racktop` from scans.

Installing Witness (Linux)

1. Retrieve a copy of the Linux `hiavd` binary by going to the web page of one of the BrickStor HA Head Nodes <https://<BrickStor Admin0 IP>:8443>. The binary comes as bzip2 compressed file.
2. Upload the compressed file to desired Linux system
3. Extract file contents using tar `tar vjxf ha-witness-linux-22.0.0.tar.bz2`
4. Copy the `hiavd` binary to `/usr/sbin` `cp ./hiavd /usr/sbin`

5. Make the binary executable `chmod 555 /usr/sbin/hiavd`
6. Configure `hiavd` as `systemd` service so that it starts and stops with the operating system.
 - a. Contact RackTop support for an example `systemd` configuration file.

TIP

For Linux systems using SE Linux feature such as RHEL or CentOS be sure to properly setup security labels for the files, directories and the service user account. If you're unsure how to do so, disable SE Linux or contact RackTop support for assistance.

TIP

Be sure to allow inbound and outbound ICMP and TCP port 4746 when the OS firewall is being used.

Forming HA Cluster

1. Using BrickStor SP Manager select one of the Head Nodes and navigate to the System tab.
2. Select Setup HA Cluster. This will bring up the HA setup wizard window.
3. In the HA wizard window fill in the appropriate information

Setup HA Cluster

General Requirements:

- All members powered on and able to ping each other via non-hb0 address.

Node Requirements:

- Connected to shared enclosure with one or more disks.
- Common pool visible but not imported by both nodes.
- Directly connected via Ethernet cable for heartbeat.
- Staged VNIC named 'hb0' created on heartbeat network interface.

Witness Requirements:

- HA service running and listening on HA comms port.
- Not member of another cluster.

Local Node		Remote Node		Witness	
10.0.0.1	X	address	10.0.0.2	X	address
192.255.0.1	X	heartbeat	192.255.0.2	X	(hb0)
.....		root pwd		

Common Resource Group Physical Interface
Interface name on nodes for HA resource group creation.

aggr0 ▼

Common HA Comms Port (advanced)
All members will use this port to communicate with each other (default 4746).

4,746

[Create/Modify](#) [Cancel](#)

Local Node - the node you are currently managing. Lets call it the first node.

Remote Node - the second Head Node.

Witness - HA witness server.

Address - IP address or a hostname.

Heartbeat - Heartbeat network interface directly connecting both Head Nodes.

root pws - Root user password for each of the HA Head Nodes.

Common Resource Group Physical Interface - sets data interface for the first Resource Group. It will also be used as a default data interface for additional Resource Groups.

Common HA Comms Port (advanced) - HA communication port. This allows changing from the default TCP port 4746.

TIP

To change the configuration of an existing HA cluster follow the same steps and enter new information. This can be handy should the IP addresses or another Default Resource Group interface needs to be established.







Managing High Availability

After the BrickStor HA cluster is formed it is managed from HA section in BrickStor SP Manager software.

The **HA Cluster** section will present dynamic action buttons that will become visible depending on cluster status.




Table 4. HA action buttons

Button	Action
	Adds new Resource Group
	Adds unmapped pool to HA configuration
	Configures advanced HA settings such as polling intervals, timeouts and failover on loss of data network
	Disables the Cluster
	Enables the Cluster. This action is only shown when HA Cluster is disabled.
	Rebalances the Cluster. Distributes Resource Groups according to their configured Preferred Node property. This action is only shown when at least one Resource Group is not on its preferred node.

Along with action buttons find a round status balloon which changes in color depending on the HA cluster health state. Green is what you expect to see when everything is healthy otherwise the color will change followed by a message like the one below. You can also hover over the status balloon for status message.

- Green - all HA components are healthy
- Orange - one or more components are degraded and HA reliability is impaired.
- Red - one or more components are faulted and HA functionality is in critical state.
- Purple - commit change is in progress.

HA Cluster Settings

Clicking the gear icon  next to **HA Cluster** allows the tuning of several advanced settings.

WARNING

Take extra care manipulating the following settings. It is highly advised to consult with RackTop support before changing the default values.

Configure HA Cluster Advanced

[Load Defaults](#)

Auto Move Resource Groups

Auto move if network link changes

Power off if unresponsive

Power off if export pool times out

Export Pool Timeout - 15.00 seconds +

Sensors

Sensor Poll Rate - 5.00 minutes +

[Launch HA Cluster Setup](#) [Apply](#) [Cancel](#)

Table 5. HA Settings

Setting	Default Value	Description
Auto move if network link changes	unchecked	enables/disables HA failover on loss of data network connectivity.
Link change delay	3 seconds	Waits <i>n</i> seconds after link state changes to down state before initiating failover.
Power off in unresponsive	checked	When enabled, healthy node will forcefully power off unresponsive node using lights-out interface in order to safely facilitate failover.
Power off if export pool times out	checked	When enabled, healthy node will forcefully power off peer node using lights-out interface in the event pool export timeout period is exceeded in order to safely facilitate failover.
Export Pool Timeout	15 seconds	On failover waits <i>n</i> seconds before forcefully powering off failed node. See Power off if export pool times out
Sensor Poll Rate	5 minutes	HA sensor polling interval in minutes


Disabling and Enabling HA Head Nodes

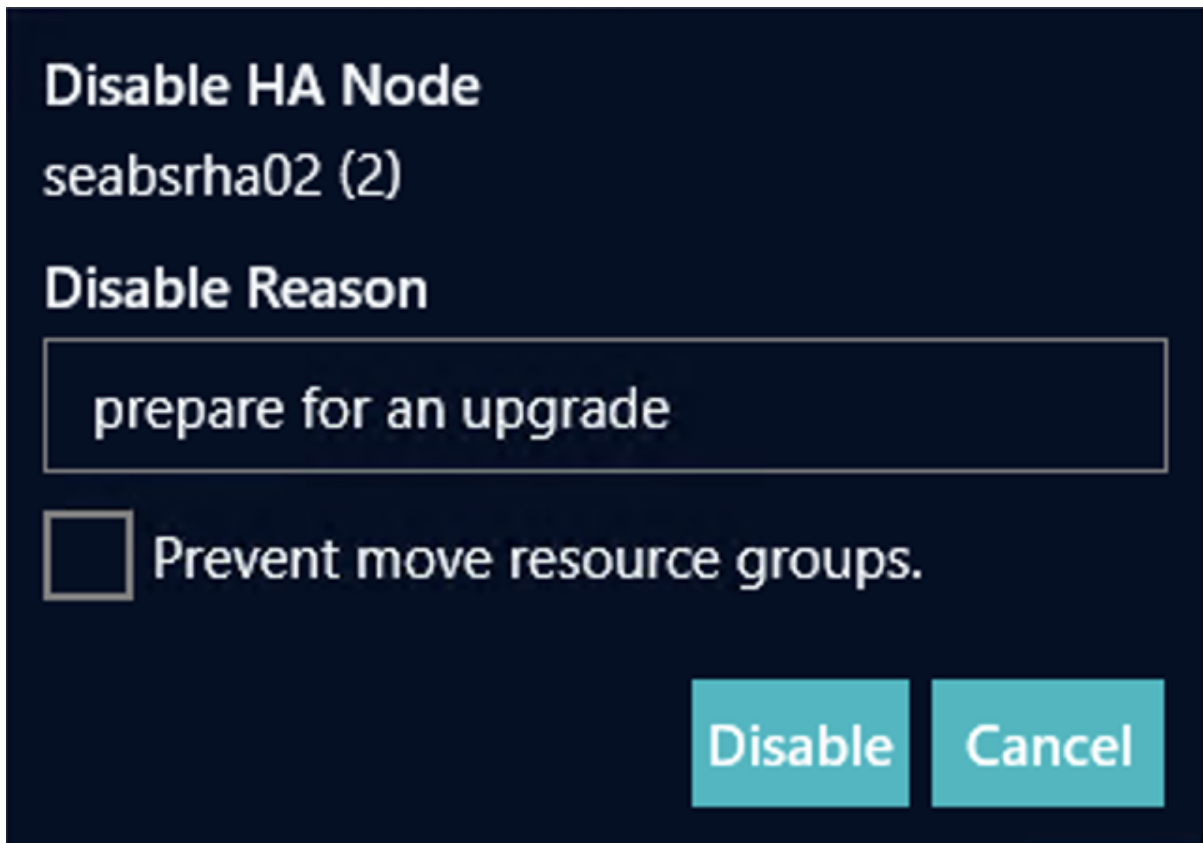
HA Cluster Head Nodes can be enabled or disabled. Disabling a node allows for taking one node out of the cluster for maintenance. When disabled the nodes do not assume any Resource Groups or participate in failover/move operations.

Enabling HA Head Node returns in back into the HA cluster. Once enabled the node can assume Resource Groups and participate in failover/move operations.

To Disable an HA Head Node:

1. Using BrickStor SP Manager connect to one of the HA Head Nodes or select it from the list.
2. In the Details pane, select the HA tab.
- 3.

Under HA Cluster mouse over desired node and click stop button . Disable HA Node dialog box will open with additional options.




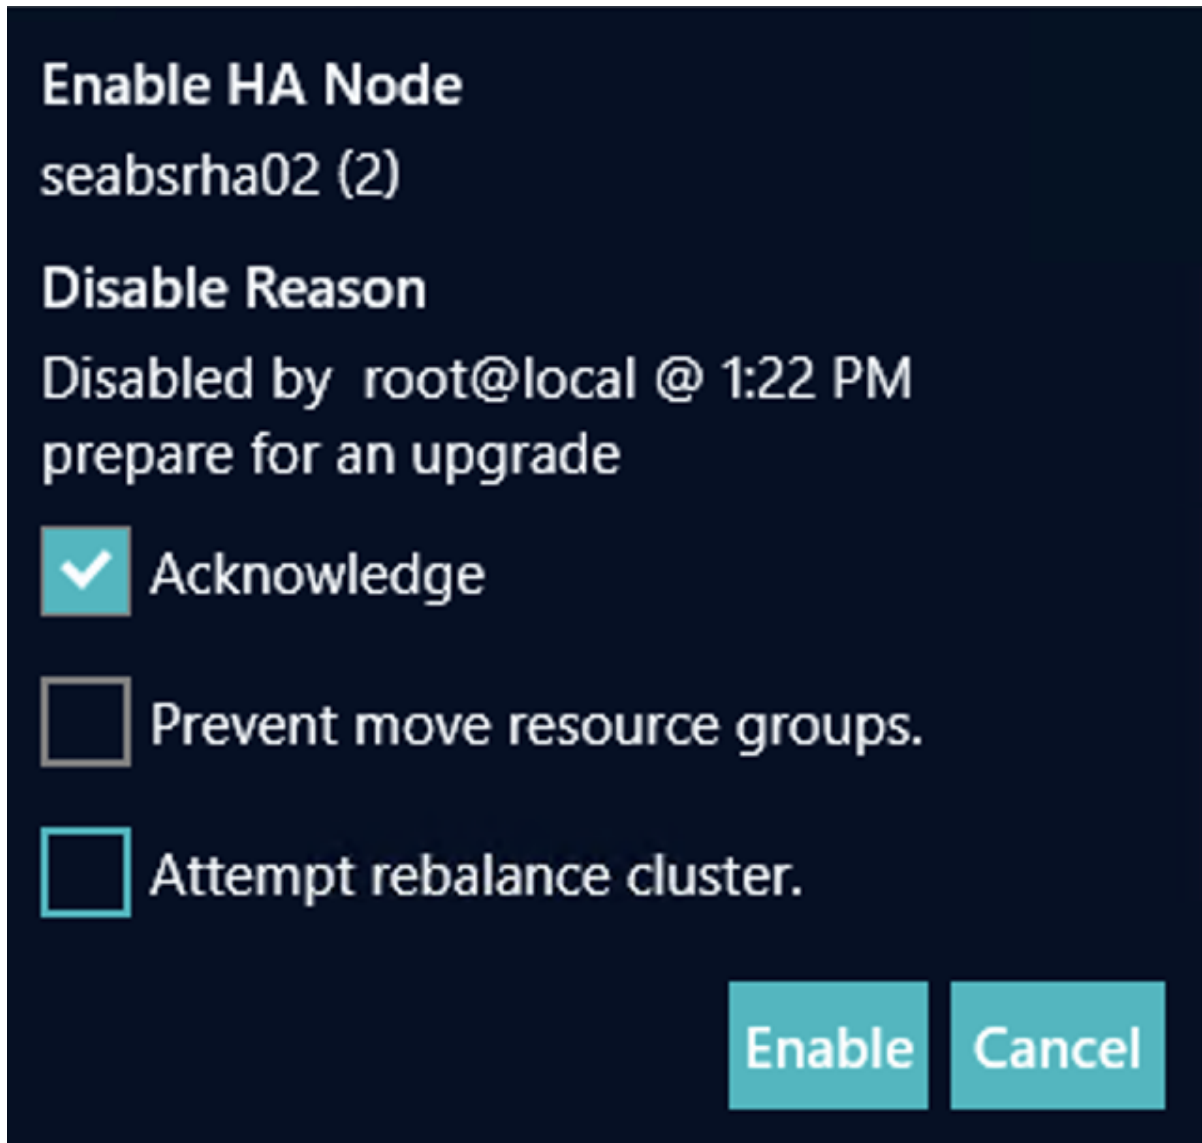
4. In the Disable HA Node dialog enter **Reason** message. This message will be used in the event log to provide a detailed explanation for this operation.
5. (optional) Check **Prevent move resource groups** to prevent automatically moving Resource Groups to the other HA Head Node. With this option selected all active Resource Groups on this node will become unavailable.
6. Click the **Disable** button.

WARNING

Avoid using **Prevent move resource groups** option. This will result in loss of data availability for all active Resource Groups on this node.

To Enable an HA Head Node:


1. Using BrickStor SP Manager connect to one of the HA Head Nodes or select it from the list.
2. In the Details pane, select the HA tab.
3. Under HA Cluster mouse over desired node and click play button . Enable HA Node dialog box will open with additional options.



4. In the Enable HA Node dialog check **Acknowledge** box to confirm delivery of the message provided during disabling operation.
5. (optional) Check **Prevent move resource groups** to prevent automatically moving Resource Groups to this HA Head Node once enabled.
6. (optional) Check **Attempt rebalance cluster** to attempt to perform rebalance operation now. This option can result in loss of data availability. First, familiarize with [Moving Resource Groups](#)

before attempting to use this option.

7. Click the **Enable** button.

TIP Enable button  will only show when an HA Head Node is disabled.

Managing Resource Groups

The initial Resource Group is created when an HA cluster is formed, however, more can be created after the fact. Additional Resource Groups only apply to systems with two or more storage pools.

When creating a Resource Group, the simple configuration contains a single pool and a single VNIC over the default interface defined during cluster creation.

In other more complex configurations it is possible to create multiple VNICs, define VLAN tags, set MTU size, choose alternate data interfaces and/or define static routes to each VNIC.

Existing Resource Groups can be managed by modifying the configured properties or by moving them manually between the HA Head Nodes.

Resource Group Properties

- **Description** - text describing the purpose of this Resource Group (ex: "User Data")
- **VNIC** - Data sharing [VNIC](#) IP address in the form of CIDR notation (ex: 192.168.0.1/24)
- **Route** - Button for adding a static route for a given VNIC
- **Pools** - [Storage Pool](#) selection
- **Select All** - Checkbox for showing/hiding [Unmapped Pools](#)
- **Node** - Node where the Resource Group currently resides. When not specified, Resource Group will become [Unmapped Resource Group](#).
- **Preferred Node** - Node where the Resource Group will reside after a Rebalance action. When set to None it will be ignored.

Unmapped Resource Group


Resource Groups that are not assigned to any HA Head Node are referred to as unmapped. Any resources allocated to this group are offline and unavailable for access. The Storage Pool(s) associated with this Resource Group will be in the exported state and VNIC configuration will not be present.



Resource Group States

There are two state messages that can be displayed next to the Resource Group name. The state messages will only show when a given Resource Group has Preferred Node property set moved either manually or automatically. Hover over the state message to display a detailed message showing an event timestamp and a reason.

The state messages can be safely ignored or remediated as needed.

- "temp" - Indicates that this Resource Group resides not on its Preferred Node. This state would show when Resource Group was manually moved to another HA Head Node by a system administrator. To remediate, click the Rebalance icon  to move Resource Groups to their preferred nodes.
- "auto-moved" - Indicates that this Resource Group has been moved to its Preferred Node by a Rebalance operation.

WARNING

Moving Resource Groups is a disruptive process and should be planned accordingly!

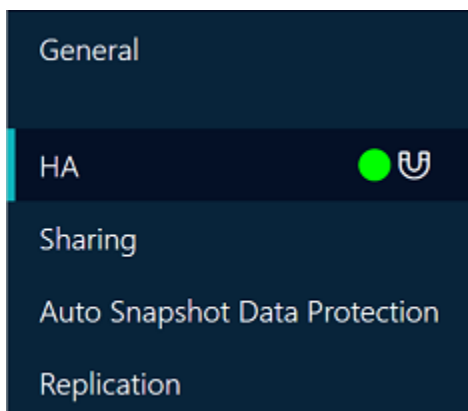
Resource Group Health


Resource Group health status is relayed via a balloon which will change colors accordingly. To see a detailed reason and timestamp of the last status change hover over the Resource Group or a status balloon.

- Green - all Resource Group components are healthy
- Orange - one or more Resource Group components are degraded and HA reliability is impaired
- Red - one or more Resource Group components are faulted and HA functionality is in critical state
- Purple - change commit is in progress
- Grey - this Resource Group is [unmapped](#)


Creating Resource Groups


1. Using BrickStor SP Manager connect to one of the HA Head Nodes or select it from the list.
2. In the Details pane, select the HA section.



3. Hover over the HA Cluster and then click the plus-R icon  to add a Resource Group. This will bring up the HA Resource Group creation dialog.



Another way to add a Resource Group is by hovering over one of the nodes and clicking the plus button .



4. In the HA Resource Group dialog, enter the required information:

Create HA Resource Group Advanced

Description

VNIC

CIDR address

VNIC address required (example 1.2.3.4/24).

Add VNIC

Pools Show All

p02 on data on habstr01 (106)

p01 on data on habstr01 (106)

Node

None - Unmapped Resource

Preferred Node

None

Create Cancel


- a. **Description** - Enter meaningful text describing the purpose of this Resource Group (ex: "User Data")
- b. **VNIC**
 - i. **CIDR address** - Enter the IP address using CIDR notation (ex: 192.168.0.1/24)
 - ii. **Route** - (optional) Add a static route for a given VNIC. Clicking this button will enter an [Advanced Resource Group creation view](#).
 - iii. **Pools**
 - A. Select at least one pool to be added to this Resource Group.
 - B. **Select All** - (optional) When checked this will show [Missing Pools](#).
- c. **Node** - Select the initial node where the Resource Group will reside once created.
- d. **Preferred Node** - (optional) Select the node where the Resource Group will reside after a Rebalance action.


- Click the **Create** button.

Creating Advanced Resource Groups


Creating advanced Resource Groups allows configuring additional properties for multiple VNICS, VLAN tags, and use interfaces other than the default cluster data interface.


- Using BrickStor SP Manager connect to one of the HA Head Nodes or select it from the list.
- In the Details pane, select the HA section.



- Hover over the HA Cluster and then click the plus-R icon  to a Add Resource Group. This will bring up the HA Resource Group creation dialog.



Another way to add a Resource Group is by hovering over one of the nodes and clicking the plus button .



- In the HA Resource Group dialog, enter the required and optional information.

Create HA Resource Group

User Data X

VNIC	Over	VID	MTU	Description
192.168.0.100/24	aggr0 (default) ▼	192	9,000 ▼	data Route
172.16.5.2/30	aggr0 (default) ▼	172	1,500 ▼	replication Route

Route → 🗑️

Add VNIC

Pools Show All

- p02
- p03 on Stealth Projects on seabsrha01 (1)
- p01 on SEDemo on seabsrha01 (1)

Node

seabsrha01 (1) ▼

Preferred Node

seabsrha01 (1) ▼

Create
Cancel

- a. Click the **Advanced** button to show advanced property fields.
- b. **Description** - Enter meaningful text describing the purpose of this Resource Group (ex: "User Data")
- c. **VNIC**
 - i. **CIDR address** - Enter IP address using CIDR notation (ex: 192.168.0.1/24)
 - ii. **Over** - Select a physical data interface for this VNIC to be created over.
 - iii. **VID** - Enter a VLAN ID.
 - iv. **MTU** - Enter a custom value for Maximum Transmission Unit (MTU). By default this will use "Auto" value to inherit MTU size of the physical interface.
 - v. **Description** - Label describing this VNIC (ex: Replication).
 - vi. **Route** - (optional) Adds a static route for a given VNIC. Multiple entries are allowed.

- A. **Destination** - Route destination using CIDR notation (ex: 0.0.0.0/0)
 - B. **Gateway** - Route gateway IP address. When VNIC IP address is already defined the value will default to the first host address of the subnet. (ex: 192.168.0.1)
 - vii. **Add VNIC** - (optional) Adds an additional VNIC.
 - viii. **Pools**
 - A. Select at least one pool to be added to this Resource Group.
 - B. **Select All** - (optional) When checked this will show [Missing Pools](#).
 - d. **Node** - Select the initial node where the Resource Group will reside once created.
 - e. **Preferred Node** - (optional) Select the node where the Resource Group will reside after a Rebalance action.
5. Click the **Create** button.

Moving Resource Groups

Resource Groups can move between HA cluster nodes automatically or can be manually triggered using BrickStor SP Manager. An automatic move can result by either a Rebalance operation or node failure.

Manual moves typically take longer compared to a failover since the HA nodes are deconstructing and reconstructing resources, whereas in a failover event the failed node is dead and we are only reconstructing. Move times can also vary depending on the system's configuration complexity. Having an unusually large amount of file systems, VNICs, static routes all contribute to extending the move/failover time. It is best to keep the configuration simple whenever possible and rather add more HA clusters to distribute complexity into multiple smaller configurations. This concept also reduces the outage impact or blast zone for the entire solution.

The Moving Resource Groups action is disruptive and should only be used during system maintenance and upgrades. It does take only several seconds and most SMB/NFS clients are designed to recover from long IO waits. However, extra care should be taken to properly plan and execute this action according to own environment.

To move a Resource Group

- | | |
|----------------|--|
| WARNING | Moving Resource Groups is a disruptive process and should be planned accordingly. |
| WARNING | Move requests do not trigger a change request and will execute upon clicking the Move button. |

1. Using BrickStor SP Manager select one of the Head Nodes and navigate to HA section.
2. Click an arrow icon next to Resource Group to be moved. This will bring up the Resource Groups move dialog.
3. Make your selections to continue or click the **Cancel** button to abort
 - a. **Selected** - Select one or more Resource Group(s) with a single operation by using the checkboxes next to them.

- b. **All on node** - All Resource Groups on the specified node. An additional node selection drop down box will show.
 - c. **All unmapped** - All unmapped Resource Groups
 - d. **All** - All Resource Groups
 - e. **To** - Destination HA Head Node where desired Resource Groups are to be moved to.
 - f. **Set preferred** - Set/change Preferred Node to destination node used.
4. Click **Move** to execute this move request

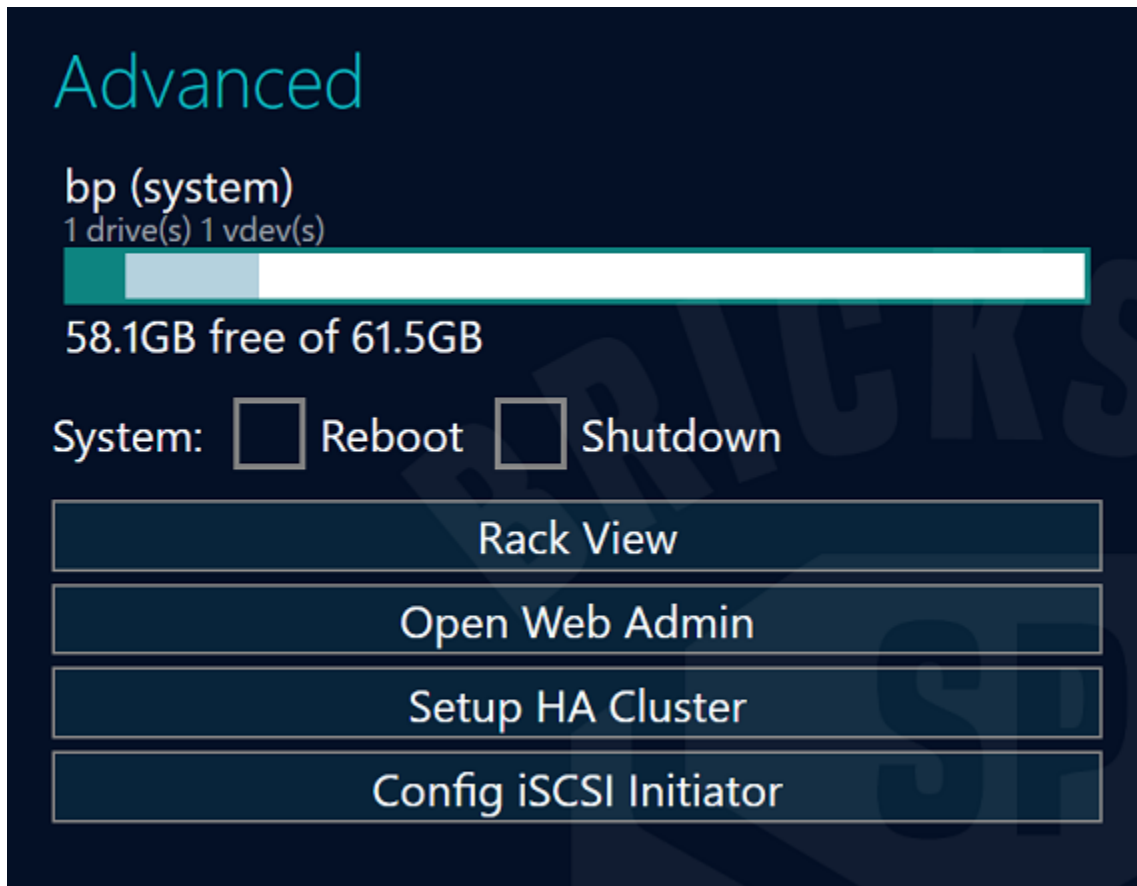
iSCSI Initiator

The iSCSI Initiator feature allows the BrickStor SP Appliance to consume block storage exported by third-party SAN solutions. Once configured, the data pools can be created using exported LUNs with BrickStor SP acting as SMB and/or NFS gateway leveraging all of the best security features. iSCSI multipathing (MPIO) is encouraged and can be configured by adding multiple iSCSI targets.

Configuring iSCSI Initiator

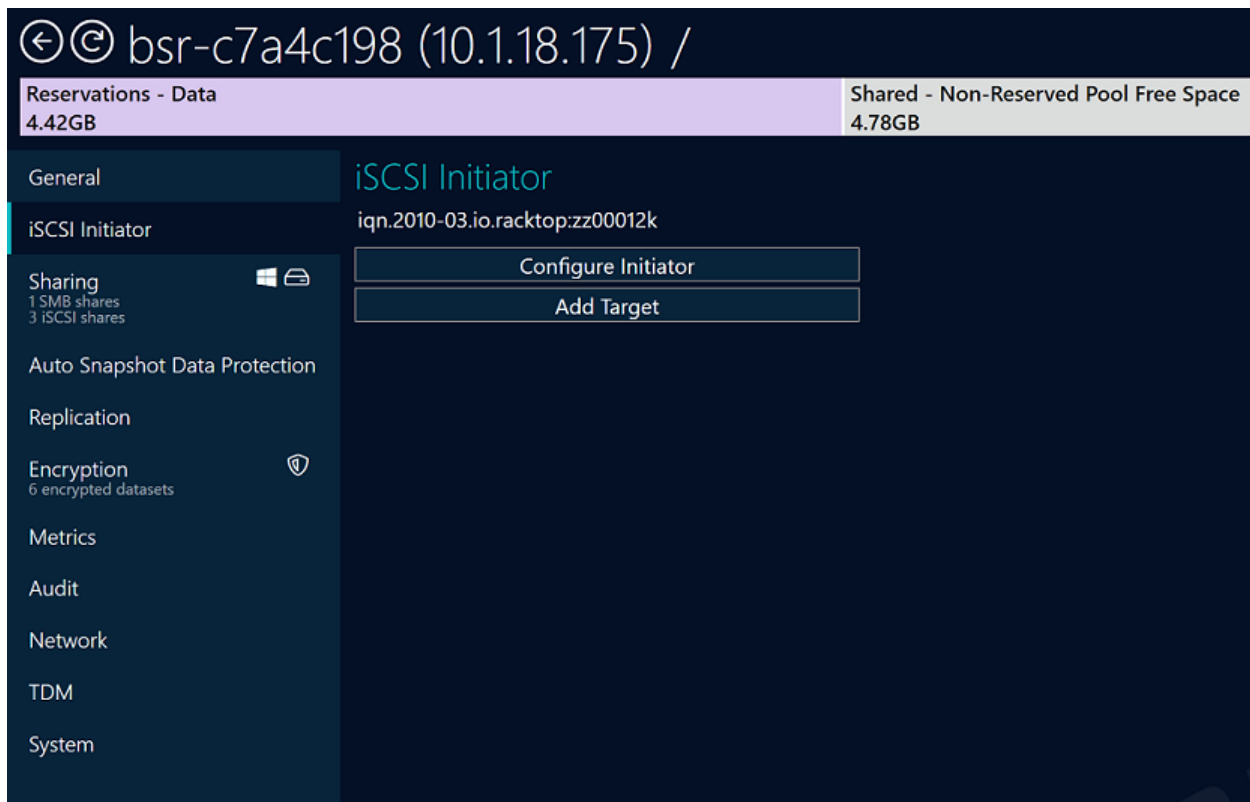
To configure the BrickStor as an iSCSI Initiator, complete the following steps:

1. In the Connections Pane, select the Appliance level.
2. In the Details Pane, click the System tab.
3. Click on the **Config iSCSI Initiator** button at the bottom right of the Details Pane.



A new tab will be created in the Details Pane labeled **iSCSI Initiator**. Navigate to this tab to see the following information:

- IQN is a unique iSCSI Qualified Name (IQN) assigned to the BrickStor 'iqn.2010-03.io.racktop:' followed by a system's serial number (i.e. 'iqn.2010-03.io.racktop:zz00012k'). This field is read-only and cannot be edited.
- A button to **Configure Initiator**
- A button to **Add Target**



Continuing to configure the iSCSI Initiator

Once the **Configure Initiator** button is selected, a pop-up window is presented where you can set optional CHAP name and secret when authentication is required by the Target. The Challenge Handshake Authentication Protocol (CHAP) enables authenticated communication between iSCSI initiators and targets. When you use CHAP authentication, you define CHAP user names and passwords on both the initiator and the Target system.

For **IQN**, this is the unique iSCSI Qualified Name (IQN) assigned to the BrickStor 'iqn.2010-03.io.racktop:' followed by a system's serial number (i.e. 'iqn.2010-03.io.racktop:zz00012k'). This field is read-only and cannot be edited. You can use this for copy/paste when adding your Initiator on the Target side.

For the Initiator's CHAP Name input the Initiator's CHAP Name expected by the Target. You can either use the already assigned IQN by selecting the associated **Use IQN** button or enter an alphanumeric free-text name.

For the Initiator's CHAP Secret input the Initiator's CHAP Secret expected by the Target (must be a value between 12-16 alphanumeric characters).

Select **Apply** to apply the changes.

IQN

iqn.2010-03.io.racktop:zz00012k

Initiator's CHAP Name

Use IQN

Initiator's CHAP Secret

Apply Cancel

Adding an iSCSI Target(s)

Once the **Add Target** button is selected, a pop-up window is presented where you can add an iSCSI Target to your Appliance. Adding an outside Target will make available to you additional volumes (LUNs) which can be added to existing pools assuming the Target has shared those volumes with you as an Initiator.

For **Initiator IQN**, the assigned IQN is again presented and is not editable. You can use this for copy/paste when adding your Initiator on the Target side.

For the **Target IQN**, you must enter the name of the desired iSCSI Target ensuring the name follows one of three formats:

- iSCSI Qualified Name (IQN) - iqn.yyyy-mm.reverse-domain-name:unique-name
- World Wide Name (WWN) - wwn.0123456789ABCDEF
- Enterprise Unique Identifier (EUI) - eui.0123456789ABCDEF

For the **TPG Tag**, a numerical value can be entered to correspond to a specific Target Portal Group (TPR) on the Target.

For **Target IP Addresses**, the IP address(es) of the desired iSCSI Target must be entered.

As enabling a two-way CHAP is optional, you can Select **Enable CHAP** when the Target requires authentication and enter a CHAP name and secret.

Select **Add Target** when all required fields are completed to apply the changes.

Initiator IQN

Target IQN

Identifier must conform to one of the following formats:
iqn.yyyy-mm.reverse-domain-name:unique-name
wwn.0123456789ABCDEF
eui.0123456789ABCDEF

TPG Tag

-0+

Target IP Addresses


One or more addresses required. One per line. Example:
1.2.3.4
1.2.3.5:999

Enable CHAP

Once added, the Target name, Target IP, connectivity status, and a representation of the connected volumes will be presented. Should the connectivity to the Target change, those same volumes will be shown as being offline.

After connecting, the configured settings to the Target can be modified using the gear icon to the right of the iSCSI Target name.

A Target can also be removed using the trash can icon to the right of the iSCSI Target name.



Compliance Reports

BrickStor SP Manager provides various exportable reports that can be accessed from the System Menu tab on the appliance level.

Compliance reports cover permissions management, data protection, data disposition reporting and other reports that are valuable for security and compliance with internal policies and government regulations. The compliance reports are designed to provide evidence of continuous compliance with standard data related controls.

Accessing Compliance Reports

To access compliance reports, complete the following steps:

1. In the Connections pane, select the appliance level.
2. Right-click and select **Open Compliance Reports**.

Select Reports by Category

When viewing a compliance report, you can select a report by category.

Favorite Reports

You can designate a report to display in favorites list by clicking the star outline.

Export Reports

You can export reports to PDF format.

Audit Log

The Audit Log displays a list of administrator actions performed through both BrickStor SP Manager and the BrickStor API. The system associates these actions with the user ID of the admin. It also displays any optional commit messages entered when the changes were committed.

Accessing the Audit Log

To access the Audit Log, complete the following steps:

1. In the Connections pane, select an appliance.
2. In the Details pane, select the Audit tab.

The screenshot shows the BrickStor SP Manager 22.0 Release Candidate (Build 1) by RACKTOP SYSTEMS. The interface is divided into several panes:

- Connections Pane:** Shows the selected appliance `bsr-d6a77c45 (10.1.29.101)` with a search bar and a list of storage pools:
 - `poolparty`: 4 drive(s) 2 vdev(s), 17.8GB free of 18.4GB
 - `global`: 16GB free of 16.5GB
 - `Data2`: 16GB free data of 16.5GB
 - `meta (system)`: 16GB free of 16GB
 - `replication`: 16GB free data of 16GB
- Details Pane:** Shows system information for `bsr-d6a77c45 (10.1.29.101)`. The **Audit** tab is selected in the navigation menu. The **Recent Actions** section displays a list of actions performed by `root@local`:
 - `unspecified commit message`
 - `AuditTdmDisableDataset` (4:04 PM (1.66ms))
 - `SnapshotDestroy` (3:57 PM (306.4ms))
 - `SnapshotDestroy` (3:57 PM (1.69s))
 - `AuditTdmModifyDataset` (3:41 PM (334ms))
 - `AuditTdmEnableDataset` (Mon 6/15 8:31 PM (0.749s))
 - `DatasetModifyReplication...` (Mon 6/15 8:31 PM (186.7ms))

3. Hover your pointer over any of the actions to display all of the API messages posted for the change.


The screenshot displays the 'Audit : Recent Actions' window in the BrickStor SP Manager. The interface is divided into several sections:

- Left Sidebar:** Shows storage metrics for 'bsr-d6a77c45 (10.1.29.101)', including 'poolparty' (17.8GB free of 18.4GB), 'global' (16GB free of 16.5GB), 'Data2' (16GB free data of 16.5GB), 'meta (system)' (16GB free of 16GB), and 'replication' (16GB free data of 16GB).
- Navigation Pane:** Contains filters for 'Scope' (bsr-d6a77c45 (10.1.29.101)), 'Search' (search path...), 'Audit Actions' (any date, now, Search, Users, Actions), 'Any Status', 'Timestamp' (Descending), and 'Related Reports' (Access Modifications, Destroy Dataset, Share Modifications, Snapshot Destroy & Release Holds).
- Main Table:** Lists recent actions with the following columns:

Timestamp	User	Action	Details	Status
6/18/2020 4:04 PM	root@local	AuditTdmDisableDataset	{"Type":"AuditTdmDisableDataset","DatasetPath":"poolpa...	Success
6/18/2020 3:57 PM	root@local	SnapshotDestroy	{"Defer":false,"Snapshot":{"CreationTimestamp":"2020-06...	Success
6/18/2020 3:57 PM	root@local	SnapshotDestroy	{"Defer":false,"Snapshot":{"CreationTimestamp":"2020-06...	Success
6/18/2020 3:41 PM	root@local	AuditTdmModifyDataset	{"Type":"AuditTdmModifyDataset","DatasetPath":"poolpa...	Success
6/15/2020 8:31 PM	root@local	AuditTdmEnableDataset	{"Type":"AuditTdmEnableDataset","DatasetPath":"poolpar...	Success
6/15/2020 8:31 PM	root@local	DatasetModifyReplicationPriority	{"Value":"Off","Inherit":false,"OldValue":"Off","OldInherit"...	Success
6/15/2020 5:14 PM	root@local	DatasetModifyPermissions	{"RecursivelyApply":false,"RecursivelyResetOwnership":"fal...	Success
6/15/2020 5:09 PM	root@local	DatasetModifyPermissions	{"RecursivelyApply":false,"RecursivelyResetOwnership":"fal...	Success
6/15/2020 5:09 PM	root@local	SmbModifyShare	{"Share":{"Name":"Data2","ReadOnlyHosts":["ReadWrite...]	Success
6/15/2020 5:09 PM	root@local	DatasetCreate	{"Dataset":{"Path":"poolparty/global/Data2","Id":"1-171876...	Success
6/15/2020 4:34 PM	root@local	DatasetDestroy	{"Dataset":{"Path":"poolparty/global/Data01","Id":"1-17187...	Success
6/15/2020 4:34 PM	root@local	SmbModifyShare	{"Share":null,"OldShare":{"Name":"Data01","ReadOnlyHos...	Success
6/12/2020 1:14 PM	root@local	AuditTdmDisableDataset	{"Type":"AuditTdmDisableDataset","DatasetPath":"poolpa...	Success
6/11/2020 4:34 PM	root@local	AuditTdmModifyDataset	{"Type":"AuditTdmModifyDataset","DatasetPath":"poolpar...	Success
6/11/2020 4:32 PM	root@local	AuditTdmEnableDataset	{"Type":"AuditTdmEnableDataset","DatasetPath":"poolpar...	Success
6/11/2020 3:43 PM	root@local	DatasetModifyReplicationPriority	{"Value":"Off","Inherit":false,"OldValue":"Off","OldInherit"...	Success
6/11/2020 11:12 AM	root@local	DatasetModifyPermissions	{"RecursivelyApply":true,"RecursivelyResetOwnership":"tru...	Success

Metrics

This tab contains various charts and graphs relating to storage capacity, cache performance, bandwidth utilization and metrics per sharing protocol.



Accessing Metrics

To access metrics, complete the following steps:

1. In BrickStor SP Manager, select the Appliance level.
2. In the Details pane, click the **Metrics** tab.

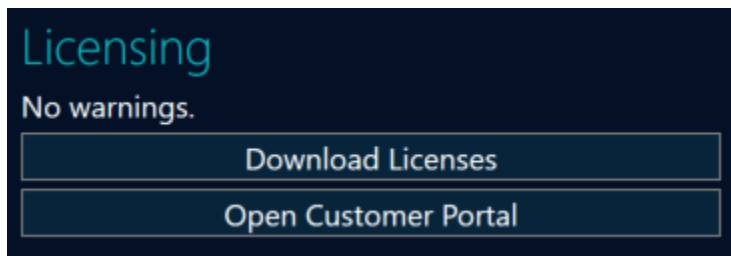
Licensing

Using the Licensing feature

The **Licensing** section displays the appliance's licensing status. It also allows for pulling down updated licenses from the MyRackTop portal.

To access licensing section select the System Tab in the Details pane using the BrickStor SP Manager.

TIP | License related **Warnings** will show here.



- **Download Licenses** allows to pull down and apply newly created or updated licenses from the MyRackTop portal to this BrickStor system.
- **Open Customer Portal** button will open web browser window to the MyRackTop Customer portal for the associated BrickStor system. The portal allows for additional management of the BrickStor(s), entitlements, and software downloads.

Command Line Reference

This reference explains the command line interface (CLI) for the BrickStor Security Platform. You can use the CLI to view, configure, and troubleshoot your BrickStor systems.

This section covers operations that can be performed via the command line. These operations are often performed via the GUI or the install tool as well.

Configuring Ethernet Address on Physical Interfaces

Network configuration model of any BrickStor appliance is essentially identical and as follows. All systems have an even number, two or four onboard 1GbE or 10GbE RJ45 interfaces clustered together to the right of the power supplies. Onboard 10GbE, 25GbE, etc. fiber or copper interfaces are typically included with every system and either reside on one of the two riser cards or installed into a special interface on the system board, on the far-right edge of the chassis. Under normal circumstances the leftmost Ethernet interface is automatically configured for the purposes of management access, which includes 'ssh' for some initial setup and for diagnostic as well as some feature configuration not currently accessible via the graphical management interface, as well as for access via the graphical management interface.

All physical interfaces are abstracted with one or multiple virtual interfaces, which do not typically share the MAC address of their underlying physical interface. Relationship of virtual interfaces to physical is many to one, meaning a single physical interface may possess one or multiple virtual interfaces. Multiple virtual interface over a single physical interface are common in VLAN tagging scenarios, which we discuss later in the document.

It is also possible to bond links into what we commonly refer to as an aggregate, more commonly known as a LAG (Link Aggregation Group), with or without LACP capability. Due to various complexities of configuration and the wide range of possible configurations we will not discuss the details of this configuration in this guide. If this is a requirement, please contact RackTop support for details.

Multiple virtual interfaces do not share a MAC address, instead each is assigned a randomly generated address, unless one is explicitly provided. We will not discuss the details of this customization here. If this is a requirement, please contact RackTop support for details.

The appliance ships with at least one already existing virtual interface named 'admin0', configured to automatically obtain an IP address via DHCP to ease initial configuration whenever possible. This is the primary management interface, meaning this is the interface used to manage the machine as an administrator, not meant for data traffic normally.

At least one data interface is required to expose files via one or more supported protocols: AFP, NFS, and SMB. For all interfaces other than management, which typically will already exist and will not need to be created or re-created, use the naming convention '**dataXX**', where '**XX**' is a non-negative numeric value starting with 0. All physical interfaces are suitable candidates, sans the first interface used for management as discussed previously. Needs vary, but a typical configuration will use 10GbE, 25GbE, and similar high bandwidth interfaces for all data access.

Virtual interfaces on BrickStorOS are configured via 'dladm', which must be created before a physical link can be used. After a system has been connected to network equipment, information about state of physical interfaces can be seen with a command in the following example. A typical output follows this general appearance:

```
# dladm show-phys
```

LINK	MEDIA	STATE	SPEED	DUPLEX	DEVICE
ixgbe0	Ethernet	down	0	unknown	ixgbe0
ixgbe1	Ethernet	down	0	unknown	ixgbe1
igb0	Ethernet	up	1000	full	igb0
igb1	Ethernet	unknown	0	half	igb1
igb2	Ethernet	unknown	0	half	igb2
igb3	Ethernet	unknown	0	half	igb3

In the above example it can be seen that the link named 'igb0' is up and configured at 1GbE. This is the physical interface on which virtual interface 'admin0' is provisioned. Typically, data interfaces will follow the naming convention prescribed earlier and use high speed interfaces, commonly identified as **ixgbeXX**, where **XX** is a non-negative numeric value starting with 0. Following is an example of establishing such a data interface over physical interface called 'ixgbe0'.

```
# dladm create-vnic -l ixgbe0 data0
```

Once a virtual interface has been created, an IP address must be assigned to this interface. IP interfaces on BrickStorOS are configured via 'ipadm'. The default 'admin0' IP interface cannot be modified since it is a temporary interface from an ipadm standpoint, instead it needs to be created persistently if a static IP address assignment is required. If you need to create a static IP, perform the following either via ssh, while connected via an IP address assigned to another interface, or directly via console of virtual console:

```
# ipadm delete-if admin0
```

```
# ipadm create-if admin0
```

```
# ipadm create-addr  $\square$ T static  $\square$ a local=x.x.x.x/24 admin0/v4
```

Where in the last command 'x.x.x.x/24' is the IP address/CIDR and 'admin0/v4' is the interface name and IP version (4 or 6). Upon creation of an IP interface, two addresses are configured, IPv4 and IPv6. For all intents and purposes IPv6 interface should be ignored usually.

VLAN Tagging

If VLAN tagging is setup on the port for trunking, you can create an interface like shown:

```
# dladm show-link
```

This will give you a list of available links for the next step, which is:

```
# dladm create-vlan -l ixgbe0 -v 10 vlan10
```

Replace ixgbe0 with an appropriate physical interface from your system and vlan10 with the name for your vlan.

NOTE | vlan name must lead with a letter and also contain at least one number.

Link Aggregation (Bonding)

If link aggregation is required, first create an aggregate and then create a vnic on top of it:

```
# dladm create-aggr -l ixgbe0 0l ixgbe1 0
```

Where '0' denotes the number that will be placed in the name 'aggr0'. After that, create a vnic on top of the aggregate:

```
dladm create-vnic 0l aggr0 data0
```

Configuring Default Gateway

If, in the previous steps, you have deleted an interface, you may not have a default gateway if the interface that was deleted was the only one on its subnet. You can find your default gateway by using:

```
# netstat -rn
```

Routing Table: IPv4

Destination	Gateway	Flags	Ref	Use	Interface
default	10.1.12.254	UG	3	5761	
10.1.12.0	10.1.12.196	U	7	2008782	admin0
127.0.0.1	127.0.0.1	UH	2	70	lo0

Routing Table: IPv6

Destination/Mask	Gateway	Flags	Ref	Use	Interface
::1	::1	UH	2	10	lo0

From there, under the flags column you are looking for a 'G', which stands for gateway and a 'default' designation under the 'Destination' column. You can add a new permanent default route using the following:

```
# route -p add default x.x.x.x
```

Where x.x.x.x is your default gateway. You can now see your default route in 'netstat -rn'

BSRAPID Configuration

By default BSRAPID is configured to listen on all interfaces on port 8443. However, the service can be configured to listen on a different port and on specific IP addresses. To configure this behavior, change the Listen Address in `/etc/racktop/bsrapid/bsrapid.conf`

Configure BSRAPID to listen on any interface with port 5443

```
ListenAddress = ".:5443"
```

Configure BSRAPID to listen only on 10.1.12.120 port 5443

```
ListenAddress = "10.1.12.120:5443"
```

Time Zone Setup

Set the time zone of BrickStor through the command line by editing the following file.

```
# tzselect
```

Then follow the prompts. A reboot is required for the changes to take effect.

NTP Setup

Preparing to Setup and Sync Time

First disable the NTP service so that you can synchronize time for the system to the NTP server. By default, the NTP service is configured to get time from the `pool.ntp.org` service.

You can enable from the command line or the GUI. To enable by command line:

```
# svcadm disable ntp
```

Next run the `'ntpdate'` command to synchronize time. This should show a current offset.

NOTE | ntp service must be disabled for ntpdate to work

```
# ntpdate <IP of Time Server>
```

If the offset was very large you can run the `ntpdate` command again to verify that clock was adjusted accordingly and offset now should be very small.

Example:

```
# ntpdate pool.ntp.org
```

```
10 Sep 08:30:08 ntpdate[7063]: step time server 129.6.15.28 offset -17971.406299 sec
```

```
# ntpdate pool.ntp.org
```

```
10 Sep 08:30:31 ntpdate[7064]: adjust time server 129.6.15.29 offset 0.002656 sec
```

NOTE

Problems with SMB authentication or AD join may be related to BrickStor's time being five minutes or more out of sync with Active Directory time.

Hosts Entries

Setting up hosts entries

Most of the time this should not be necessary, but in the exceptional cases where host name resolution is required and cannot be accomplished via DNS, static entries may be added to allow for local resolution. This activity is accomplished via `bsradm` as follows, where `'192.168.0.1'` is the address and `'othernode'` is name resolving to this address:

```
# bsradm hosts add --ip 192.168.0.1 --names othernode
```

NOTE

This may be a required step if DNS is not setup and you are connecting to an NFS datastore from ESXi.

RMM (Remote Terminal) IP Address

Your BrickStor storage appliance comes equipped with a Remote Management Module frequently abbreviated to RMM. RackTop recommends connecting this Ethernet interface as well as the `'admin0'` management Ethernet interface to a dedicated management network, if one is available. Separation of management and administration concerns from data access is a recommended best practice. This enables you to access the appliance as if you were standing in front of it with a crash cart or KVM, even when services such as SSH are down. You can use RMM to power cycle the machine or see and use the console. If RMM is already configured, you can find the IP address with this command from the terminal:

```
# bsradm hw rmm
```

```
IpSource: DHCP Address
```

```
IpAddress: 192.168.0.101
```

```
SubnetMask: 255.255.255.0
```

```
MacAddress: 00:1e:67:50:c7:c1
```

```
SnmpCommunityString: public
```

```
DefaultGateway: 192.168.0.1
```

```
Vlan: 0
```

Once you have the IP address, you can login to RMM via your browser. You will need to use Java to access the console and this will most likely require adding a security exception for the IP address in the Java control panel.

Creating Local Accounts

As root you can create local accounts that can be used for controlled access to shares as well as providing access to administrative functions such as the ability to manage BrickStor with the BrickStor SP Manager.

```
#useradd <username>
```

To set the user's password:

```
#passwd <username>`
```

Add Local Accounts to Bsradmins Group

To allow a given local account administrative access of a BrickStor appliance via BrickStor SP Manager, this account must be in the 'bsradmins' group of the appliance. To add a user to the group, run the following command, replacing username placeholder with actual local account name:

```
# usermod -G bsradmins <username>
```

Adding and removing e-mail addresses from Report Notification List

To add e-mail addresses to receive notifications from the BrickStor appliance, use the following command format at the terminal:

```
# bsradm notify add <email address> all
```

Other options besides the "all" notifications options are:

```
--system Add to system notification list
```

```
--reports Add to reports notification list
```

```
--faults Add to faults notification list
```

To list users and their notification types, use:

```
# bsradm notify show
```

And to remove users from their notification, use:

```
# bsradm notify remove <email address> --all
```

Joining Active Directory

The first step for making a CIFS share available for users is to join Active Directory, which requires several configuration steps before joining the domain will be possible. A machine account will be created for a BrickStor upon successful domain join operation. This machine account will enable users to passthrough authenticate and be either permitted or denied access to shares without

requiring separate authentication against the BrickStor. In other words, once users are logged into Active Directory, their authentication information is stored on their system and in Active Directory, and no further authentication prompts are necessary in order to access shares on a domain-joined BrickStor.

Active Directory requires certain attributes of name resolution, which usually means the BrickStor must be configured to resolve names against domain in the given instance of Active Directory to which it will be bound. BrickStor's domain setting must also be set to name of domain being joined.

First, validate what is currently configured, because no change may be necessary. Check currently configured domain with the following command:

```
# bsradm dns domain get
```

If the value reported is correct, that is, it matches the Active Directory domain name, no change is necessary. If, however a modification is necessary, change should be made with the following command, replacing placeholder 'domain.tld' with actual fully qualified Active Directory domain name:

```
# bsradm dns domain set <domain.tld>
```

Next, confirm that correct DNS resolvers are configured, and if not, make necessary changes. In most environments at least two DNS servers will be configured and BrickStor must point to these resolvers, which in typical Active Directory configurations will be domain controllers also, or commonly member servers with a dedicated DNS function.

First, validate what is currently configured, because no change may be necessary. Check currently configured domain name resolution servers with the following command:

```
# bsradm dns ns show
```

If values reported are correct, no further resolver changes should be necessary. If, however DNS servers need changing, use the following commands to add/remove entries, replacing placeholder 'address' with IP address of system being added or removed.

```
# bsradm dns ns add <address>
```

```
# bsradm dns ns remove <address>
```

NOTE

NTP must be correctly configured with accurate synchronized timing with the Domain Controller before you can join the Domain Successfully.

The command for joining the storage appliance to the domain is:

```
# smbadm join -y -u <Administrator Account> <domain.tld>
```

Where 'Administrator' is the name of the user you want to use to join the domain. This account is only used to create the computer object and does not need to be a service account. You will be prompted for a password. If the join fails, please double check your username and password and the settings in /etc/resolv.conf.

To view the domain type:

```
# smbadm list
```

Also verify that forward and reverse lookups are correct within active directory for the BrickStor.

Performance Monitoring

There are several scripts included with BrickStor for monitoring the performance of the storage portion of the system. Dtrace and kstat are powerful tools for analyzing the storage performance and behavior.

Iostat

Iostat is one of the most common tools to assess disk performance. Running `!iostat -xn 5` from the command line will result in an output similar to the following.

```

                `extended device statistics`
`r/s`    `w/s`    `kr/s`    `kw/s` `wait` `actv` `wsvc_t` `asvc_t` `%w` `%b`
`device`
`2.7`    `142.4`    `1.0`    879.8  0.6  0.0    4.0    0.1  14  2 c2t0d0
`2.0`    `1.3`    `11.3`    0.0  0.0  0.0    0.0    0.0  0  0
c0t5000C5002E4A47DAd0
`2.0`    `1.3`    `11.3`    0.0  0.0  0.0    0.0    0.0  0  0
c0t5000C5002E4B60D8d0
`2.0`    `1.3`    `11.3`    0.0  0.0  0.0    0.0    0.0  0  0
c0t5000C5002E46D2C5d0
`2.0`    `1.3`    `11.3`    0.0  0.0  0.0    0.0    0.0  0  0
c0t5000C5002E4AC53Cd0
`2.0`    `1.3`    `11.3`    0.0  0.0  0.0    0.0    0.0  0  0 c2t1d0
`2.0`    `1.3`    `11.3`    0.0  0.0  0.0    0.0    0.0  0  0
c0t5000C5002E4B0940d0
`2.0`    `1.3`    `11.3`    0.0  0.0  0.0    0.0    0.0  0  0
c0t5000C5002E4665F8d0
`2.0`    `24.3`    `11.3`    137.9  0.0  0.2    0.0    6.7  0  8
c10t50000393C8C93AF6d0
`2.0`    `24.7`    `11.3`    `137.9  0.0  0.2    0.0    6.5  0  8
c10t50000393C8C91A5Ad0
`0.0`    `0.0`    `0.0`    `0.0  0.0  0.0    0.0    0.0  0  0
c10t50000393C8C918E6d0
`2.0`    `6.0`    `11.3`    `109.4  0.0  0.0    0.0    0.1  0  0
c0t5001517BB2863DF8d0
`2.0`    `1.3`    `11.3`    `0.0  0.0  0.0    0.0    0.0  0  0
c0t5001517BB27C697Cd0
`2.0`    `36.3`    `11.3`    `269.7  0.0  0.5    0.0    13.6  0  18
c10t50000393C8C9184Ed0
`2.0`    `37.3`    `11.3`    `269.7  0.0  0.4    0.0    9.0  0  11
c10t50000393C8C93BAEd0

```


Key values to look for are %w, %b, asvc_t. It is normal for these values to increment beyond zero in a heavily loaded system, but they should usually be in relation to the overall system load. If the system is not heavily utilized, and these values are consistently high on a single device, it may indicate that the device is experiencing a hardware issue.

Zpool iostat

Zpool iostat shows extended details about a ZFS pool. Running `!zpool iostat -v 3` from the command line will result in an output similar to the following:

```
capacity operations bandwidth latency pool alloc free read write read write read write
-----
poolA 1.18T 6.98T 0 0 0 0.00 0.00 mirror
403G 2.33T 0 0 0 0.00 0.00 c0t5000C5002E46D2C5d0 - - 0 0 0 0 0.00 0.00
c0t5000C5002E4A47DAAd0 - - 0 0 0 0 0.00 0.00 mirror 403G 2.33T 0 0 0 0 0.00 0.00
c0t5000C5002E4AC53Cd0 - - 0 0 0 0 0.00 0.00 c0t5000C5002E4B60D8d0 - - 0 0 0 0 0.00 0.00
mirror 403G 2.33T 0 0 0 0 0.00 0.00 c0t5000C5002E4B0940d0 - - 0 0 0 0 0.00 0.00
c0t5000C5002E4665F8d0 - - 0 0 0 0 0.00 0.00 c2t1d0 784K 29.7G 0 0 0 0 0.00 0.00
-----
syspool 16.1G 216G 0 2 0 11.9K 0.00 0.22
c2t0d0s0 16.1G 216G 0 2 0 11.9K 0.00 0.22 -----
vmpool01 160G 3.47T 0 131 0 633K 0.00 27.46 mirror 130G 1.69T 0 19 0 68.3K 0.00 24.90
c10t50000393C8C91A5Ad0 - - 0 13 0 69.6K 0.00 10.56 c10t50000393C8C93AF6d0 - - 0 13 0 69.6K
0.00 9.15 c0t5001517BB2863DF8d0 2.89M 22.2G 0 9 0 82.0K 0.00 0.11 mirror 30.3G 1.78T 0 36 0
159K 0.00 36.40 c10t50000393C8C93BAEd0 - - 0 20 0 160K 0.00 14.72
c10t50000393C8C9184Ed0 - - 0 19 0 160K 0.00 21.07 cache - - - - - c0t5001517BB27C697Cd0
99.1G 12.7G 0 0 0 0 0.00 0.00 -----
```

In this instance, you would be looking for latency values > 100 for extended periods of time as an indicator of an overloaded system.