BrickStor SP User Guide

Release 23.3

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. Without limiting the rights of this copyright, no part of the Documentation may be modified, 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 tortuous 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.

Table of Contents

Getting Started with BrickStor SP . . . 6
- BrickStor SP Initial Configuration . . . 6
- Installing BrickStor SP . . . 6
- The BrickStor SP Console . . . 6
- Registration . . . 11
- Using BrickStor SP Manager . . . 17
- BrickStor SP Manager Overview . . . 19
- General User Layout and Conventions . . . 19
- The Rack View Interface . . . 24
- General Appliance Information . . . 27
- Appliance Sharing Information . . . 28
- Network Information . . . 29
- System Information . . . 30
- Joining Active Directory . . . 35
- Managing BrickStor SP with AD Users . . . 39

Data Protection . . . 43
- Data Protection Information . . . 43

Pools . . . 45
- Boot Pools . . . 45
- Hybrid Pools . . . 45
- RAID Performance . . . 47
- Pool Hierarchy and Containers . . . 49
- Pool Types . . . 49
- Creating Pools . . . 50
- Viewing Pools . . . 53
- Managing Pools . . . 54
- Pool Storage Utilization . . . 65
- Pool Performance . . . 66
- Pool Sharing Information . . . 68
- Pool Settings . . . 68
- Destroying Pools . . . 68

Datasets . . . 70
- Shares . . . 70
- Creating Datasets . . . 72
- Working with Datasets . . . 74
- Dataset Permissions . . . 74
- Quotas and Reservations . . . 78
- Dataset Bars . . . 79
- Dataset Storage Utilization . . . 80
- iSCSI . . . 81

Snapshots . . . 84
- Snapshot Indexing . . . 84

Restoring a file from a Snapshot . . . 85
Snapshot Holds . . . 86
Rolling Snapshots . . . 86
Clones . . . 88
Replication . . . 89
Replication Best Practices . . . 89
Understanding Peers . . . 89
Configuring a Peer Relationship . . . 90
Understanding Peer Status . . . 92
Data Protection Replication . . . 93
Data Protection Policy Configurations . . . 93
Data Replication Priorities . . . 95
Configure the Data Protection Policy for a Storage Profile . . . 95
Managing Replication Details . . . 95
Replication Transfer History . . . 96
Auto Snapshot Data Protection . . . 97
User Behavior Auditing and Analysis . . . 101
Enabling User Behavior . . . 101
User Behavior Audit . . . 102
Active Defense . . . 107
Security Incident Display and Workflows . . . 107
Assessors and Rules . . . 112
Ransomware & Malware Protection . . . 114
Insider Threat . . . 114
Threat Level . . . 120
File Recovery . . . 121
Bulk Recovery . . . 125
High Availability . . . 131
High Availability Components . . . 131
HA Cluster Architecture . . . 133
HA Scenarios . . . 134
High Availability (HA) Best Practices . . . 136
Configuring High Availability . . . 137
Prerequisites . . . 137
Setting up Witness Server . . . 138
Distributed Configuration Database (confd) Windows Install . . . 143
Forming HA Cluster . . . 147
Managing High Availability . . . 149
HA Cluster Settings . . . 149
Disabling and Enabling HA Head Nodes . . . 151
Managing Resource Groups . . . 153
Encryption and Key Management . . . 161
Managing Encryption . . . 161
Encryption Best Practices . . . 162

Self Encrypting Drives . . . 164
Drive Enrollment . . . 164
Other Self Encrypting Drive Operations . . . 165
Exporting and Backing Up Keys . . . 165
Cryptographically Erasing SEDs . . . 166
SED Protection on the Main Pane . . . 167
Transparent Data Movement (TDM) . . . 168
File Chunking . . . 168
Demand Cache . . . 168
Logical Segmentation – Enclave Elimination . . . 168
Configuring TDM . . . 168
TDM Status and Data Distribution . . . 171
Reconfiguring TDM . . . 173
Disabling TDM . . . 173
iSCSI Initiator . . . 175
Configuring the iSCSI Initiator . . . 175
Configuring Initiator authentication . . . 176
Connecting to the iSCSI Target . . . 177
Compliance Reports . . . 180
Accessing Compliance Reports . . . 180
Select Reports by Category . . . 180
Favorite Reports . . . 180
Export Reports . . . 180
Audit Log . . . 181
Accessing the Audit Log . . . 181
Metrics . . . 183
Accessing Metrics . . . 183
Licensing . . . 184
Using the Licensing feature . . . 184
Health . . . 186
Accessing the Health Tab . . . 186
Health Tasks . . . 187
Webhooks . . . 189
Microsoft Teams Webhook Connector . . . 189
Pager Duty Notifications . . . 190
Pushover Notifications . . . 191
Slack Webhook Connector . . . 192
RackTop Webhook Format . . . 193
Managing Webhooks . . . 194
Upgrading . . . 195
Pre-upgrade Considerations . . . 195
Upgrading a Single Node BrickStor using the latest BrickStor SP Manager . . . 196
Post-Upgrade Tasks . . . 198
Upgrading Distributed Configuration Database (confd) from 23.0.6 to 23.2 . . . 198

Addendum 200
Open Network Port Requirements 200

Getting Started with BrickStor SP

RackTop's Cyberstorage software, BrickStor SP, is a secure-software defined NAS platform for unstructured data. Users deploy data-centric zero trust architecture to ensure compliance and stop cyberthreats in real-time.

BrickStor SP Initial Configuration

The following instructions will guide initial configuration of the BrickStor SP Console on first install.

Installing BrickStor SP

The following steps will outline the process of installing the BrickStor SP Manager to a system.

TIP If the BrickStor SP Manager is being installed over a Virtual Machine Environment, navigate to Installing the BrickStor SP on a Virtual Machine

The BrickStor SP Console

NOTE While navigating the BrickStor SP Console, use **Tab** to navigate while editing fields. Use **Arrow Keys** to navigate to different fields on the Main Menu.

- Upon first boot, the BrickStor SP Console will launch.
- A message will direct the setting of the root password before using the appliance.

Setting a Root Password

The following instructions will guide the creation of the Root Password.

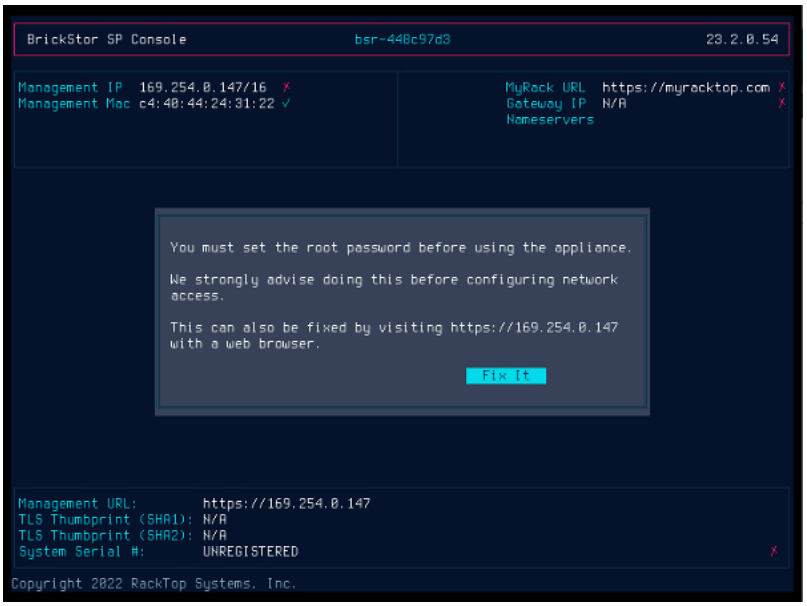

- While at the Status Screen, press **Enter**.

- The Change Root Password screen will present.

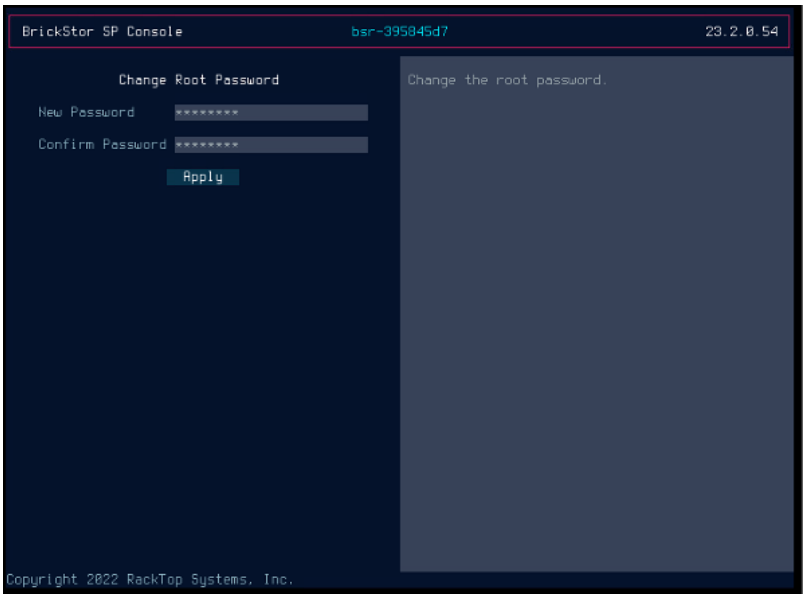

- In the supplied fields, enter a new Root Password.
- Press **Tab**.
- Enter the same Root Password in the **Confirm Password** field to confirm this change.
- Press **Tab**
- Press **Enter** to apply changes, and set the new Root Password.

- A message denoting a successful Root Password change will present, press **Enter** to proceed.

Accepting Terms and Conditions

- The BrickStor SP console will now prompt for Terms & Conditions Acceptance

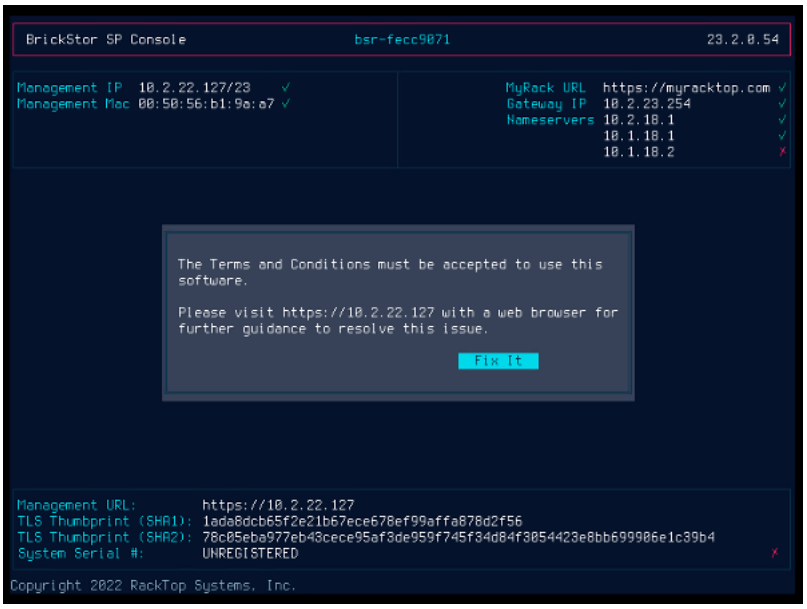

- While on the BrickStor SP Console Status screen, ensuring the 'Fix It' prompt is present, press **Enter**.

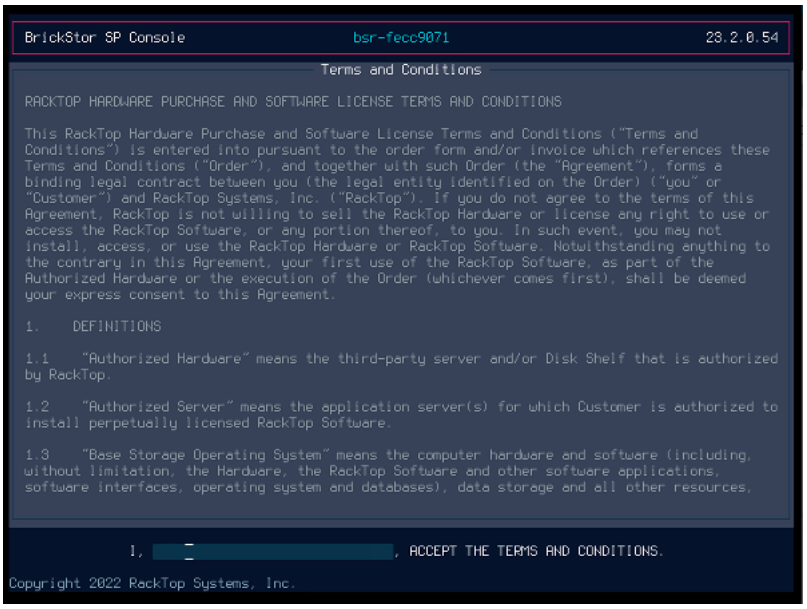

- The Terms and Conditions Screen will present.
- Read through the Terms and Conditions of the BrickStor SP License.
- When finished, press **Enter**.
- Enter user name in the supplied field.
- Press **Enter**.

Configuring Management Network

NOTE	If your appliance has obtained a network address using DHCP, you can skip this section and proceed to Registration.

- In the BrickStor SP console, A message will present noting the system as not registered.

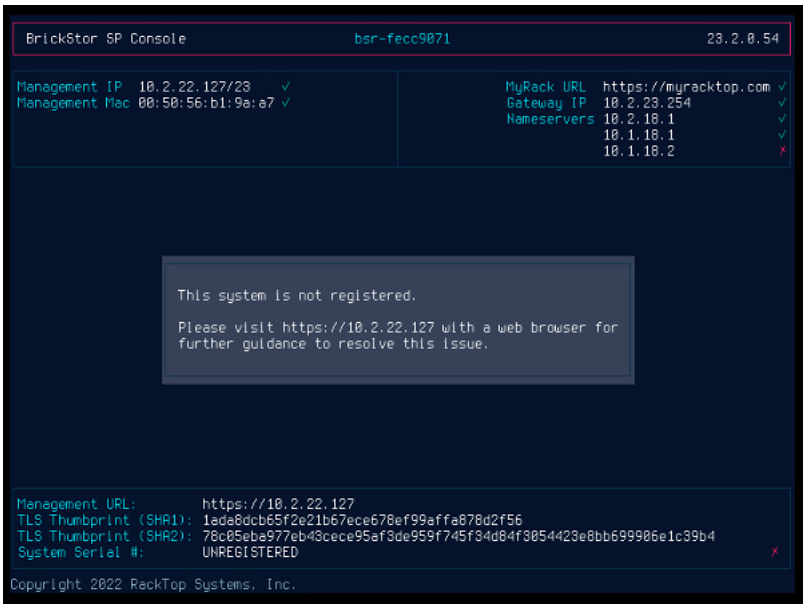

- Press **F3**.
- The following menu will present.

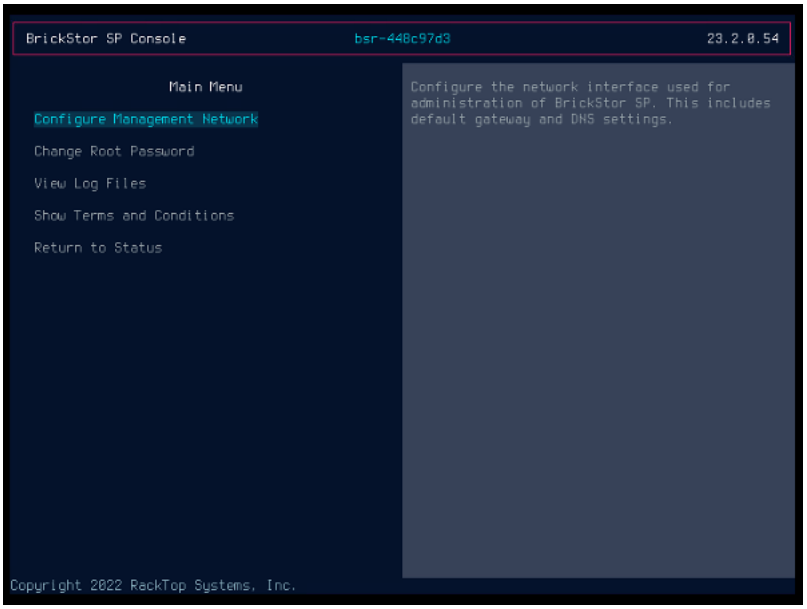

- Ensuring **Configure Management Network** is highlighted, press **Enter**.
- The Configure Management Network screen will present.

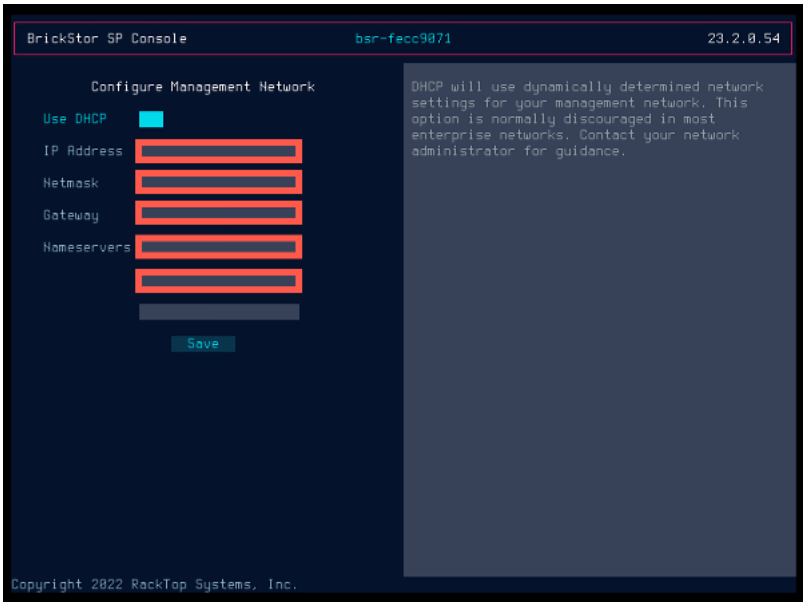

- Enter the following information:
  - IP Address
  - Netmask
  - Gateway
  - Nameservers (It is recommended to add a minimum of two)
- Press **Tab** to select **Save**.
- Press **Enter**.

NOTE

Alternatively, it is possible to utilize DHCP to Configure Management Network.

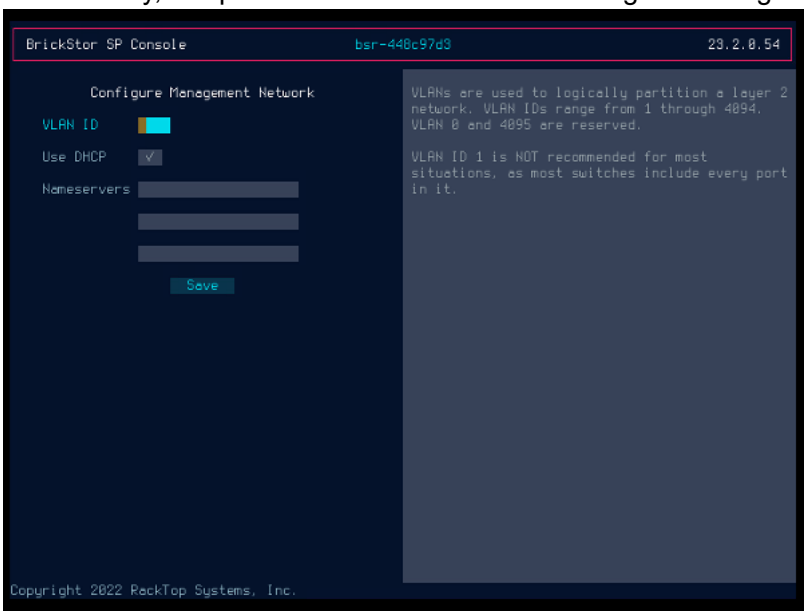

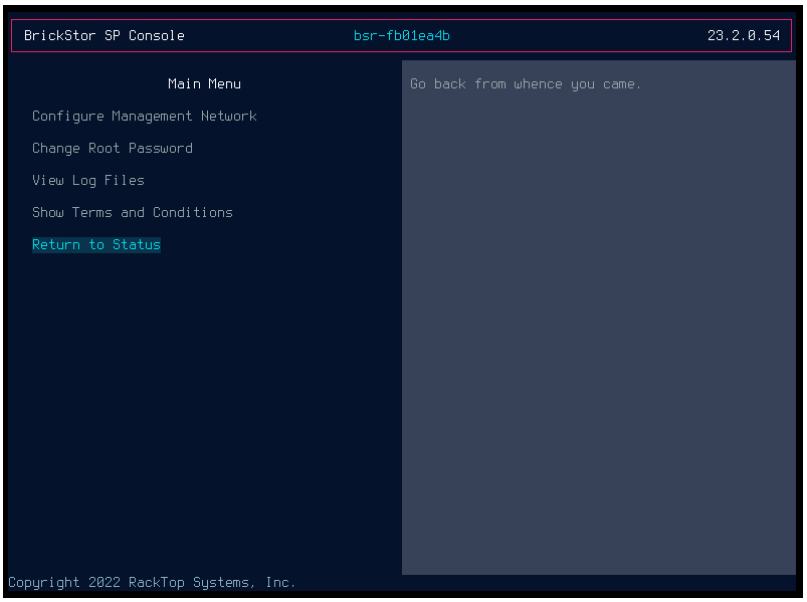

- Press **Tab** to navigate to **Return to Status**.
- Press **Enter**.
- Once finished, the BrickStor SP Console Initial Configuration has been completed.

Registration

Online Registration

In order to register BrickStor SP online, a license entitlement is needed for each installation. An entitlement is obtained by purchasing BrickStor SP software license. To view current entitlement login to myRackTop portal https://www.myracktop.com and navigate to the **Entitlements** page.

NOTE	Network must be configured for online registration.

To register a new BrickStor SP online complete the following steps:

Boot the New BrickStor SP

- After the initial boot of a new system the console will display the URL to register the appliance. Navigate to that URL to continue with the registration process.

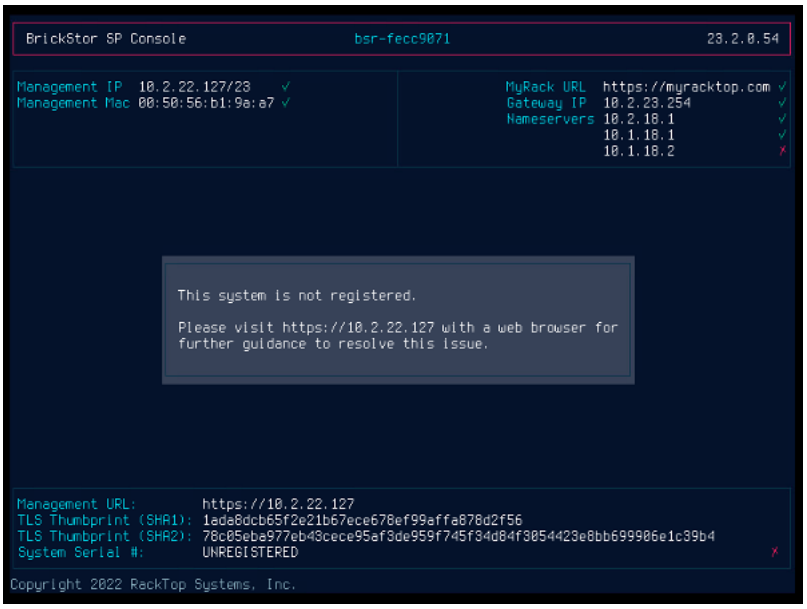

- Enter customer ID, then click **Register Now**.

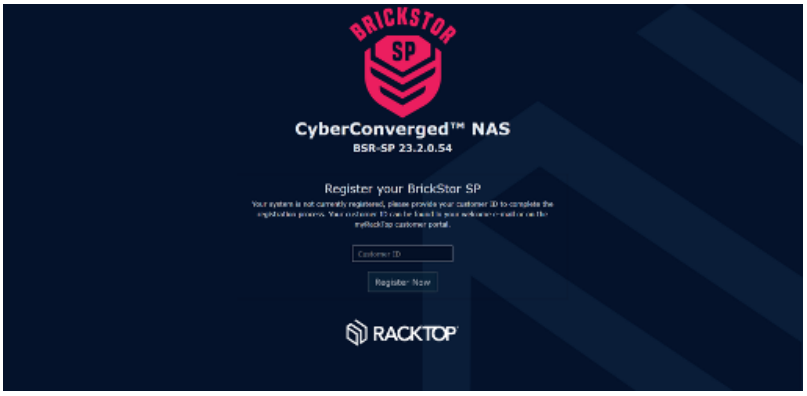

- The registration URL will present. Click the provided link.

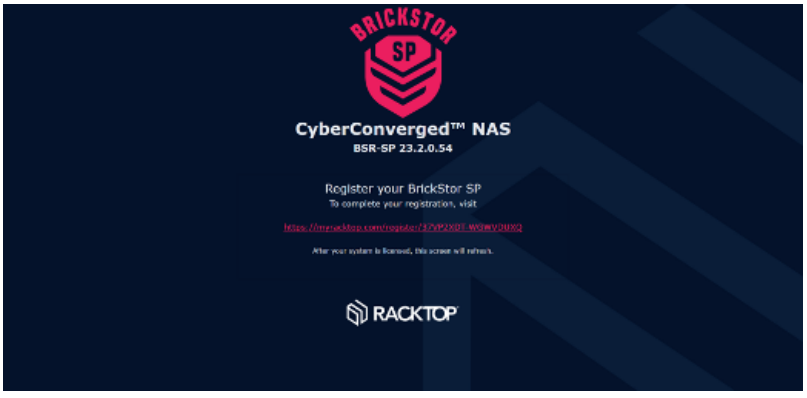

- Follow the on-screen instructions to be taken to the list of entitlements.

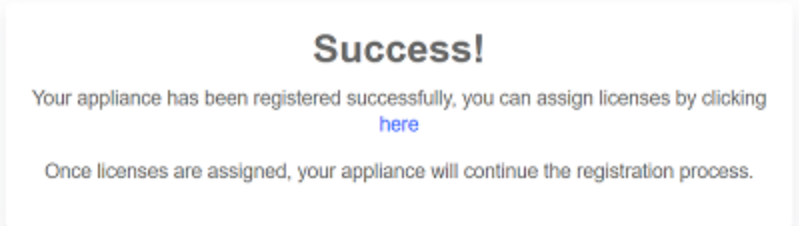

- Select an entitlement from the list and click **License Now**.

NOTE Once an entitlement is assigned to an appliance, it cannot be reused unless the appliance is deleted from the portal.

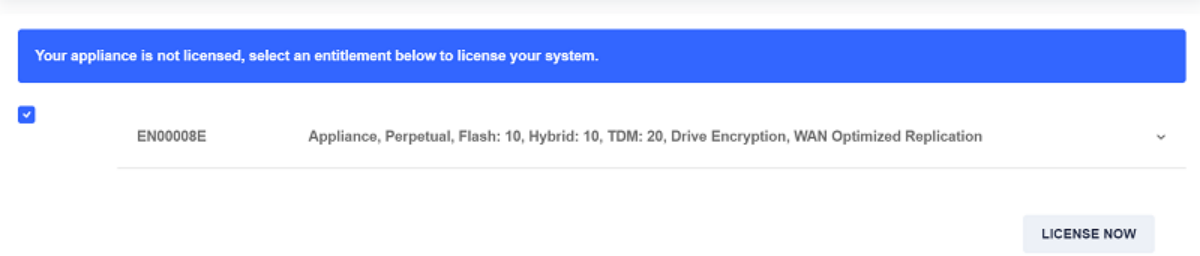

- The licenses are applied to the Brickstor SP and it is registered.

NOTE

The following will show on the BrickStor SP Console during Portal registration.

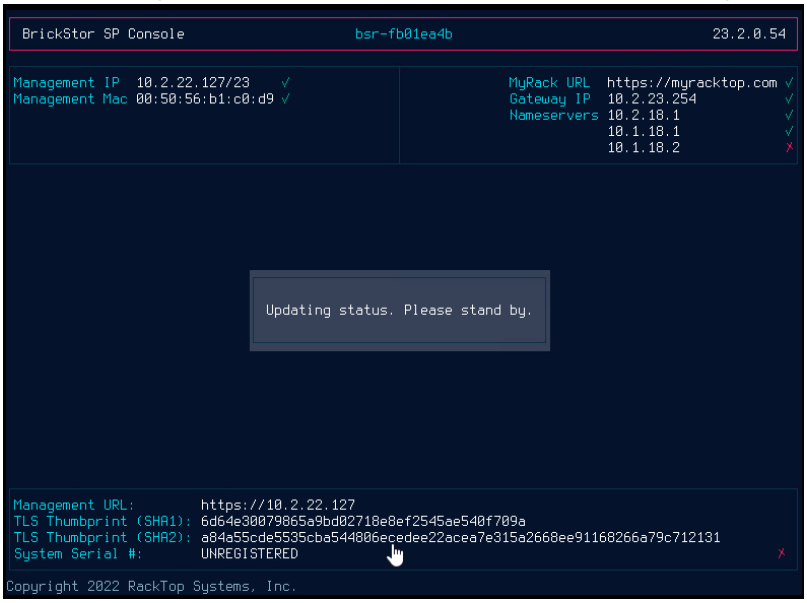

- The console of the BrickStor SP will now present a prompt.

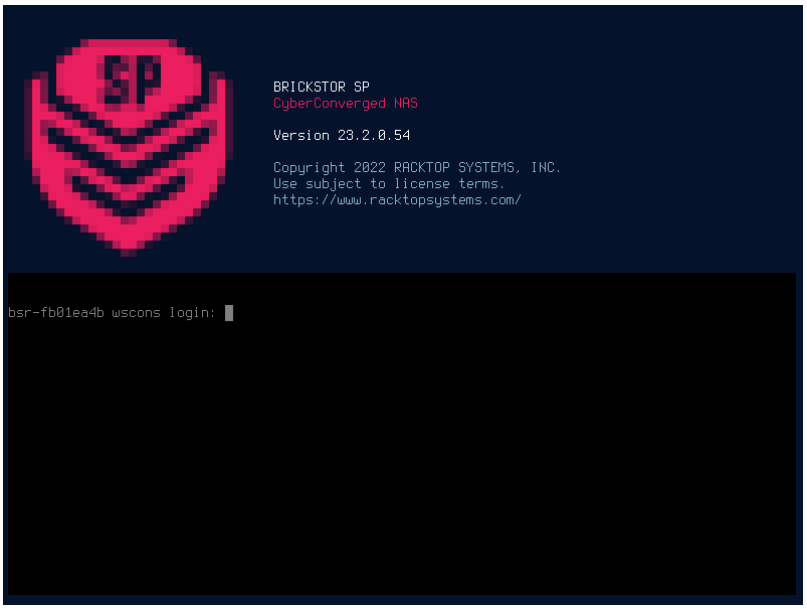

Offline Registration

In order to register BrickStor SP offline, a license entitlement is needed for each installation. An entitlement is obtained by purchasing BrickStor SP software license. To view current entitlement login to MyRackTop portal https://www.myracktop.com and navigate to the **Entitlements** page.

NOTE	Offline Registration is used only when the BrickStor SP device is not connected to the internet.

Boot the New BrickStor

1. Follow the URL and enter credentials to login.

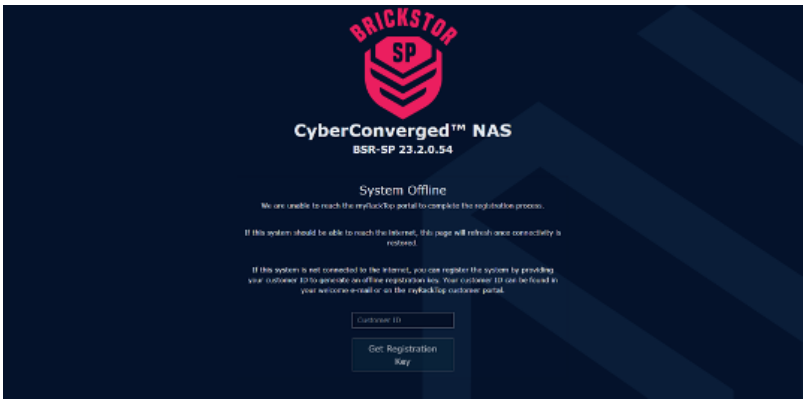

1. After entering customer ID, An offline registration key will present.

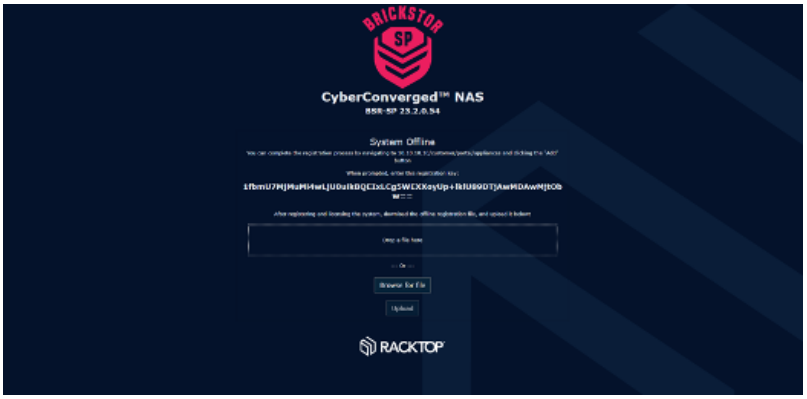

1. Copy and paste this key into the **Add Offline Appliance** option on https://www.myracktop.com. The **Add Offline Appliance** button can be found in the top right of the screen, on the **Appliances** tab.

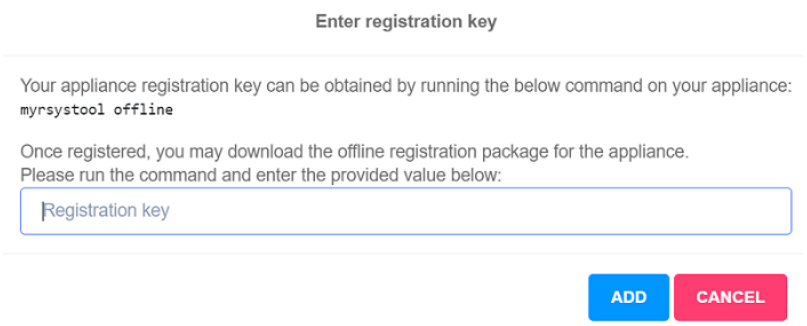

1. Click **Add** to register the BrickStor SP.
2. The BrickStor SP is now registered and ready for an entitlement to be assigned. A list of entitlements will present and can select which one you wish to assign to that system and click **License Now**.

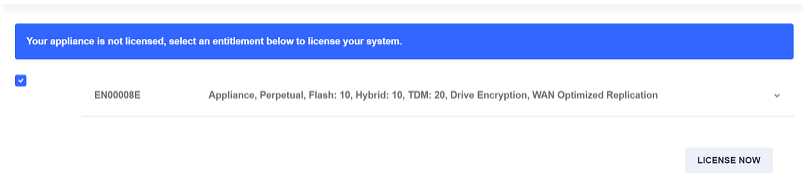

1. The license details screen will present. At the top of this list there will be a **Download Registration** button, click that to obtain the offline registration file.

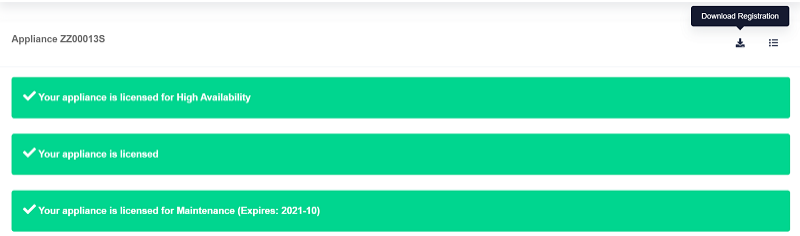

1. Take this file and upload it on the main BrickStor SP homepage. All licenses and certificates will automatically be applied to the system.

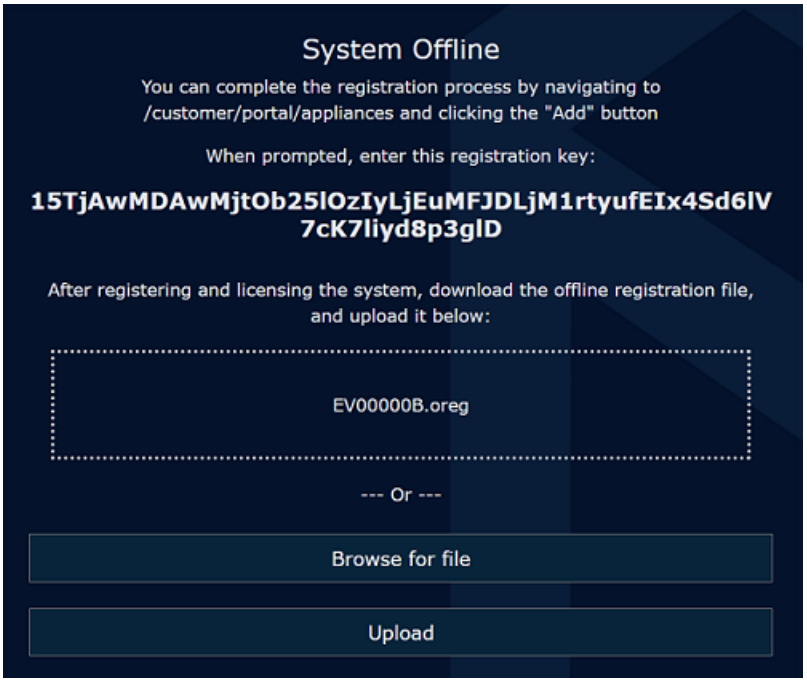

1. The BrickStor SP homepage will refresh automatically and prompt the creation of a new password.

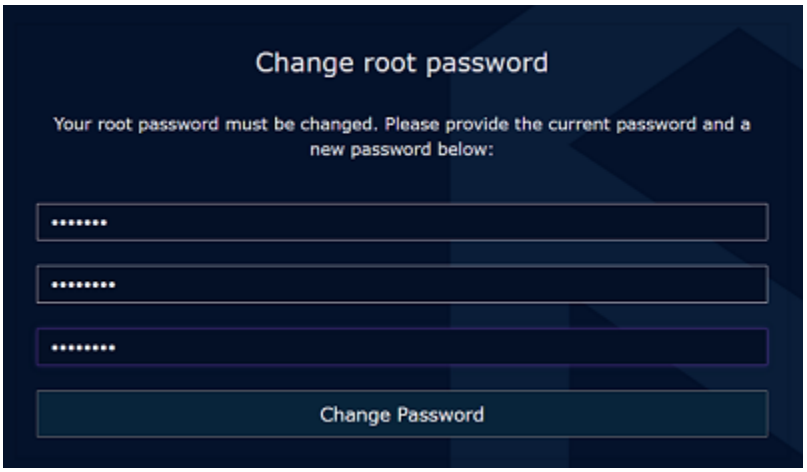

1. After creating a new password, a prompt will direct a fresh login to the BrickStor SP homepage.

1. The console of the BrickStor SP will now be at the login prompt.

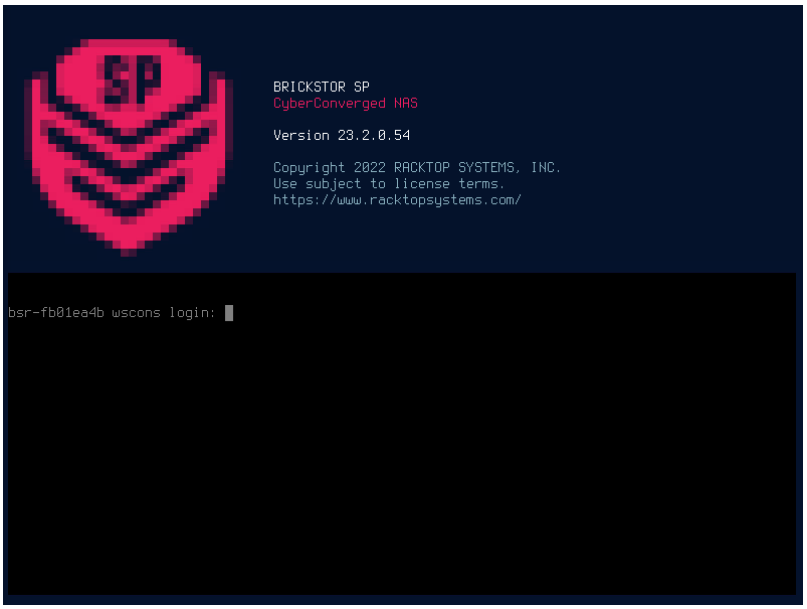

1. Return to Installing BrickStor SP.

Using BrickStor SP Manager

BrickStor SP has a user interface entitled BrickStor SP Manager that can be used to perform administrative, management, analysis, and auditing tasks.

BrickStor SP Manager can manage multiple BrickStor SP appliances. BrickStor SP Manager runs on Microsoft Windows.

To download and use 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 a BrickStor SP appliance via BrickStor SP Manager:

1. Run brickstorspmgr.exe by **double clicking** it.

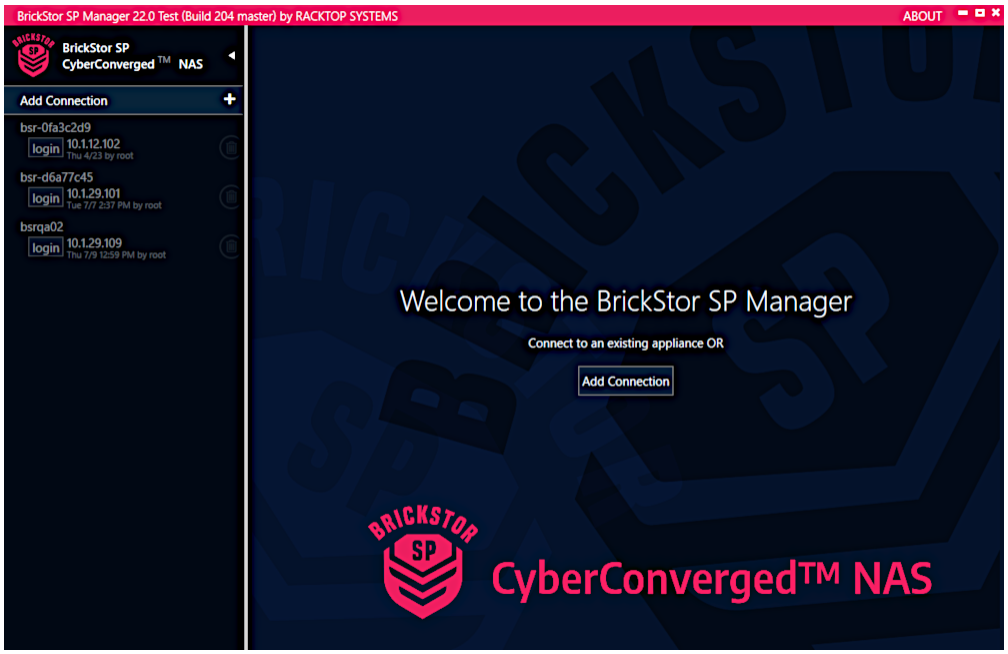

2. Click **Add Connection** to add a connection to a BrickStor SP appliance.

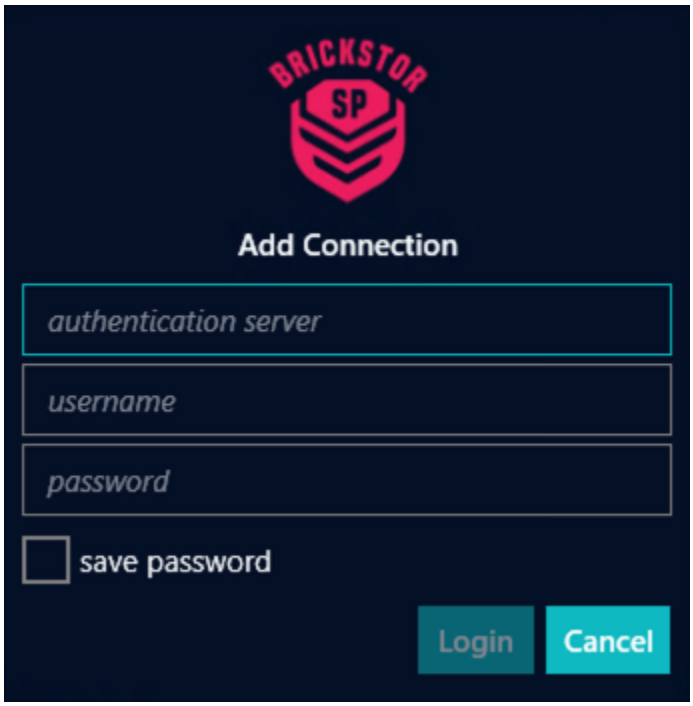

3. In the Add Connection dialog box, enter the following:

- 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 Overview

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 for management of the BrickStor SP at a granular level.

The BrickStor SP Manager is capable of managing a single BrickStor SP 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:

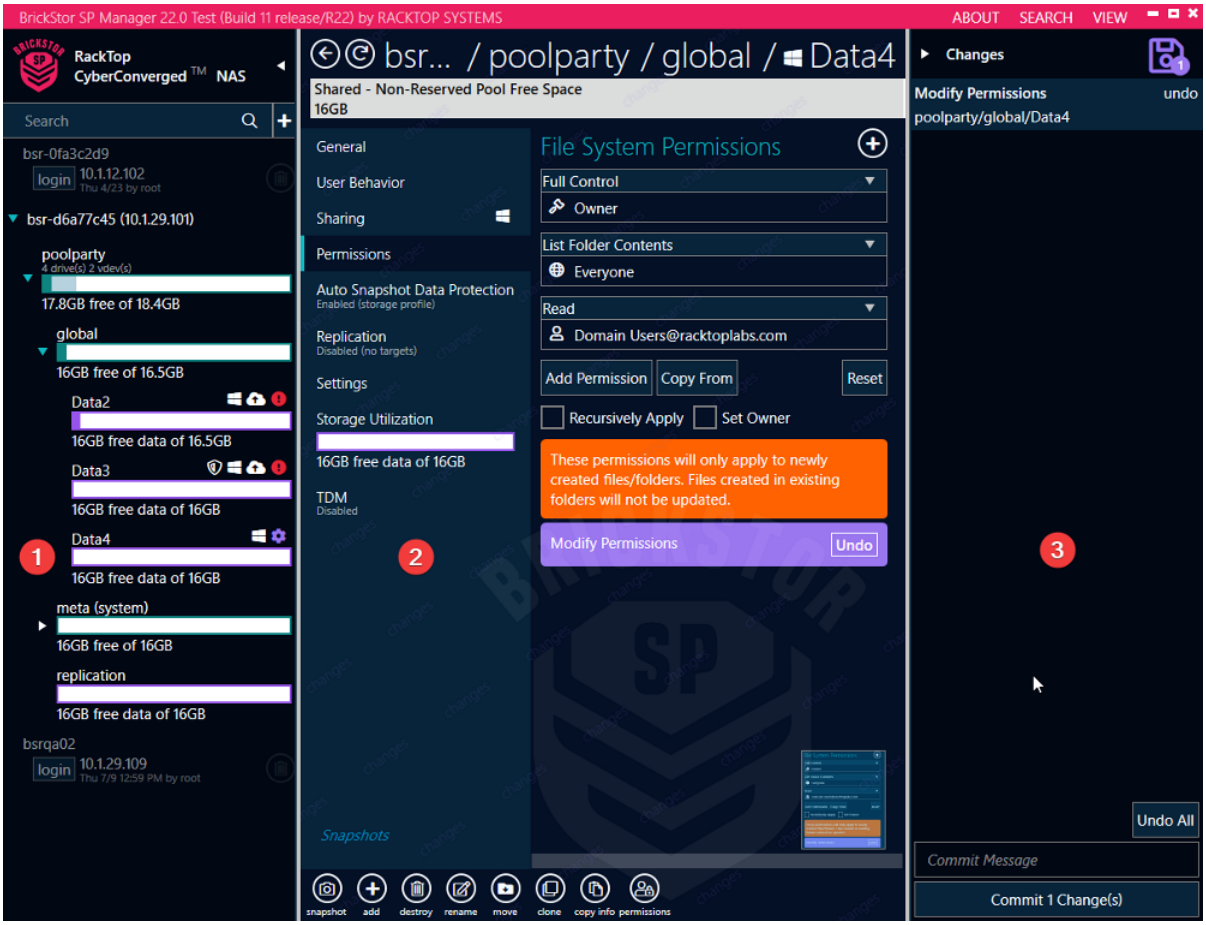

1. the Connections pane
2. the Details pane
3. the Changes pane

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.

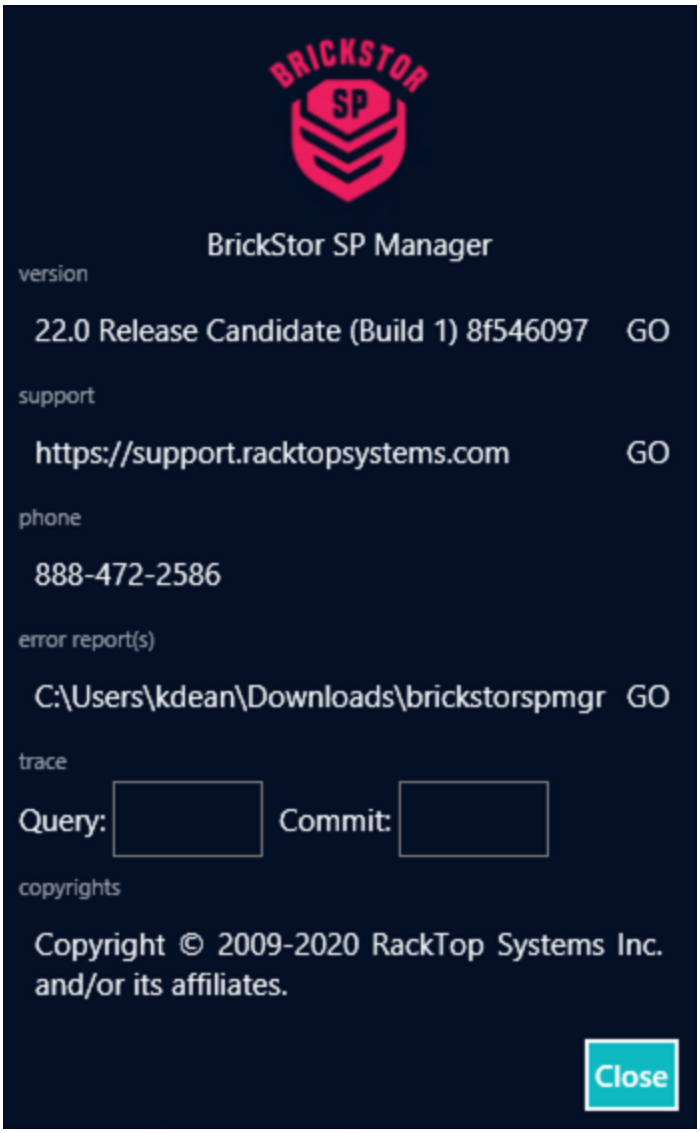

TIP By setting a value, for example 5GB, in the Trace Query and Commit box will create a local log on the machine running BrickStor SP Manager with all of the GUI requests and responses.

Search Menu

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

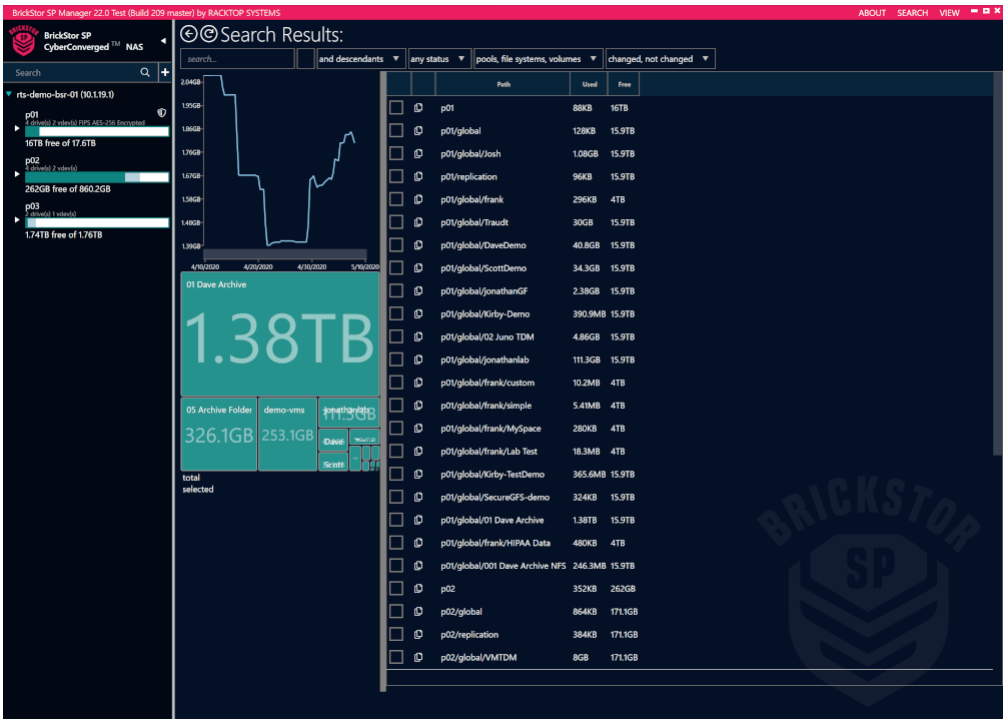

View Menu

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.

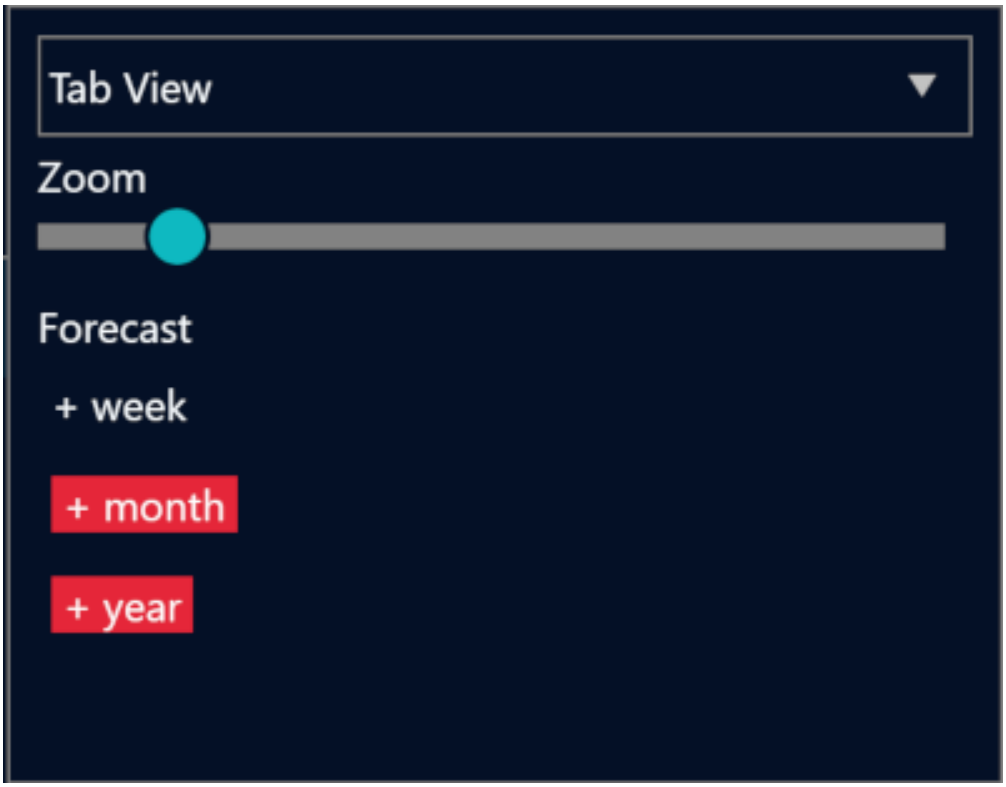

Finally, you can adjust Zoom properties, which change the width of columns in all views.

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

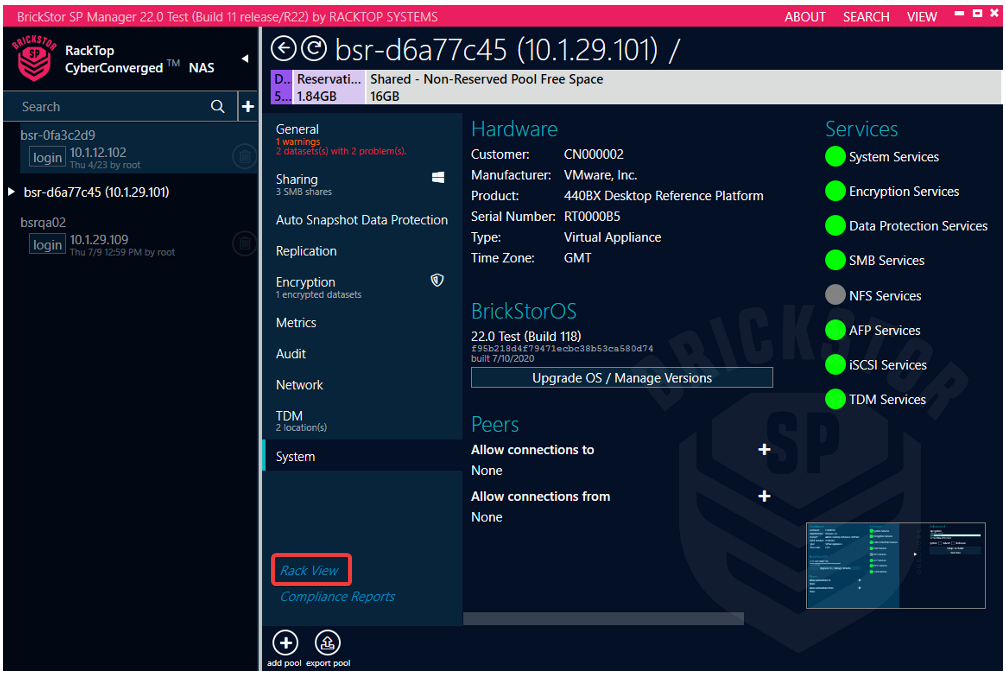

You can use Rack View to easily view and modify your appliance hardware. Rack View allows users to add or modify pools and vdevs and gives visuals that allow users to see what changes will occur to the system's hardware prior to committing them. It will also display various diagnostic information such 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.

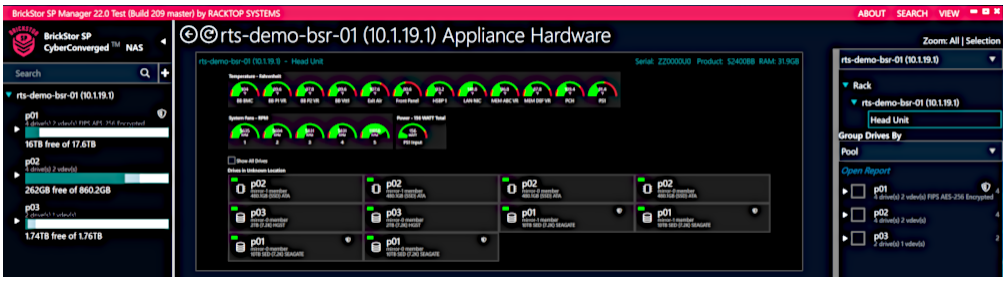

The right-hand side of Rack View also allows you to group the drives in the appliances based on certain properties such as pool, make, and vdev type. To change the grouping type, select the dropdown under Group Drives By and then select how you want to group them. When hovering over one of these groups, 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.

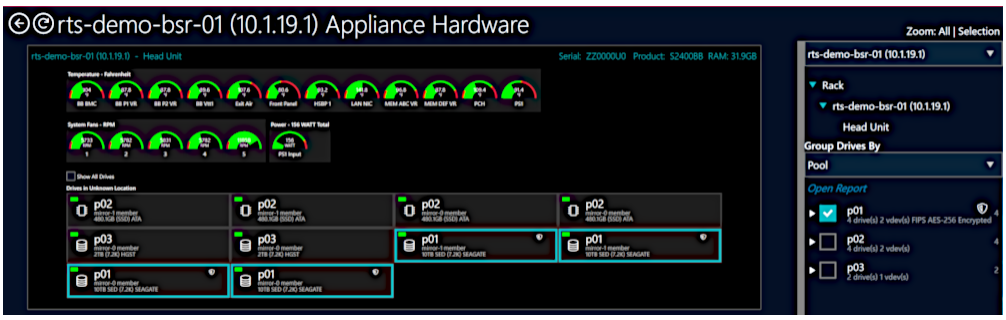

Accessing Rack View

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 in it 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.

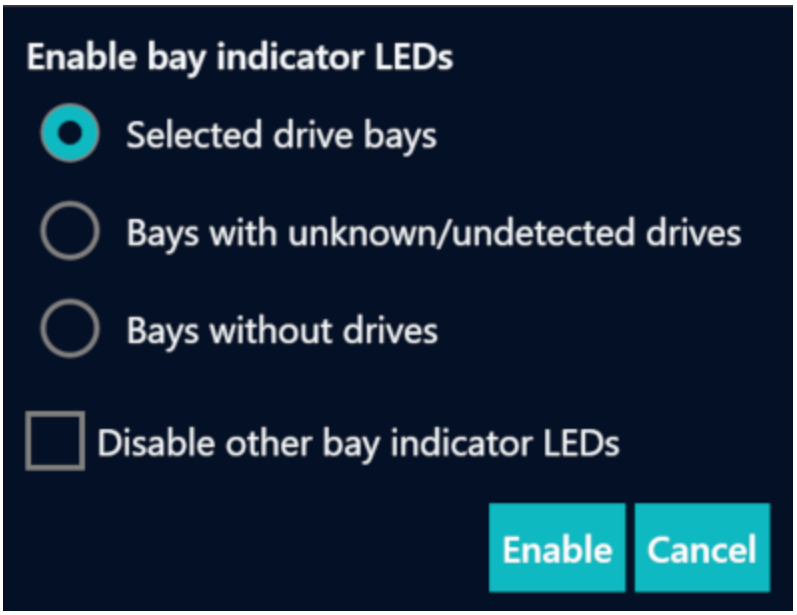

Drives with their indicating LEDs enabled will have a blinking orange indicator on Rack View as well as on the physical drive on the appliance.

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.

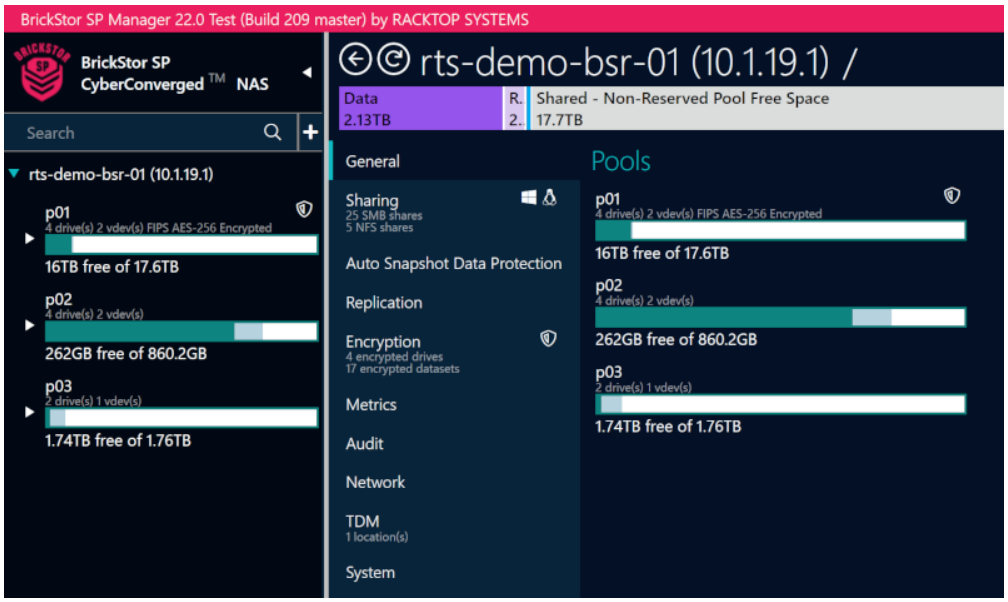

Appliance Sharing Information

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.

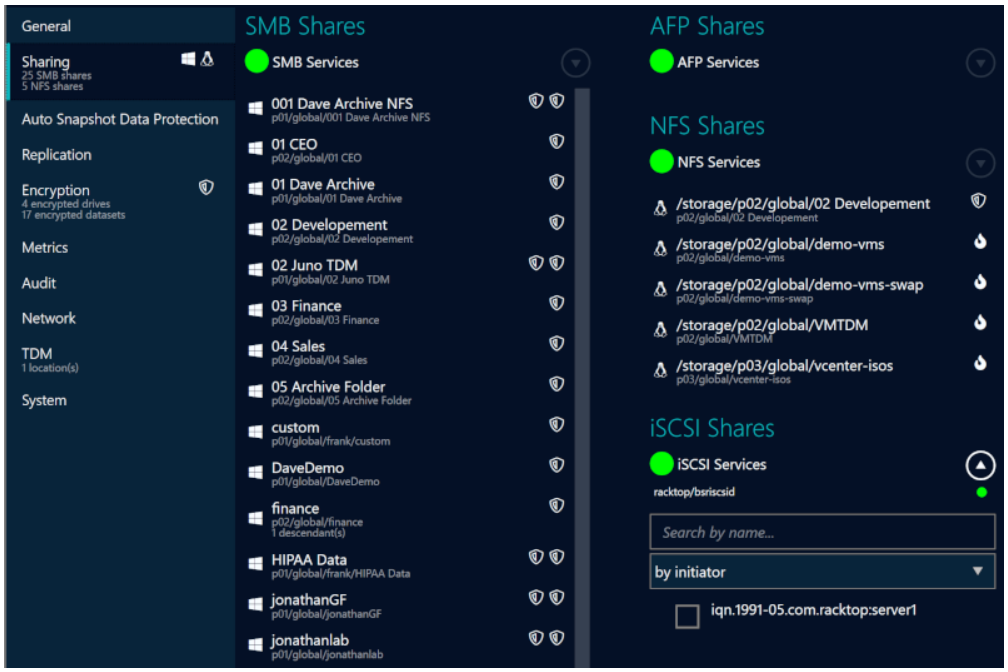

Network Information

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.

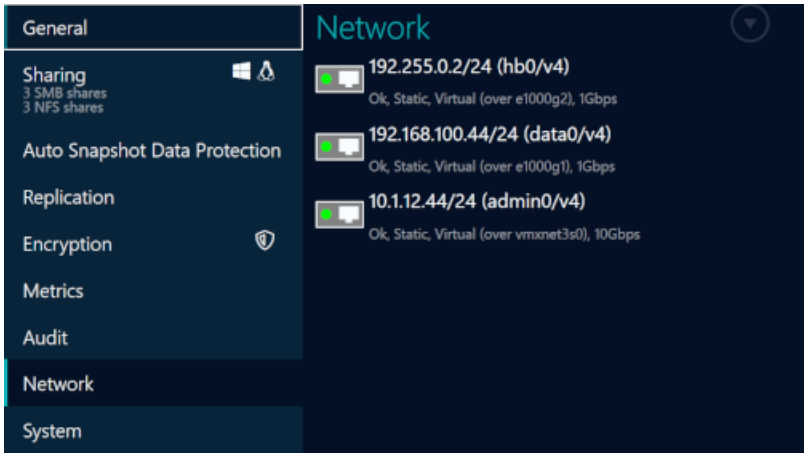

System Information

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 admins 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.

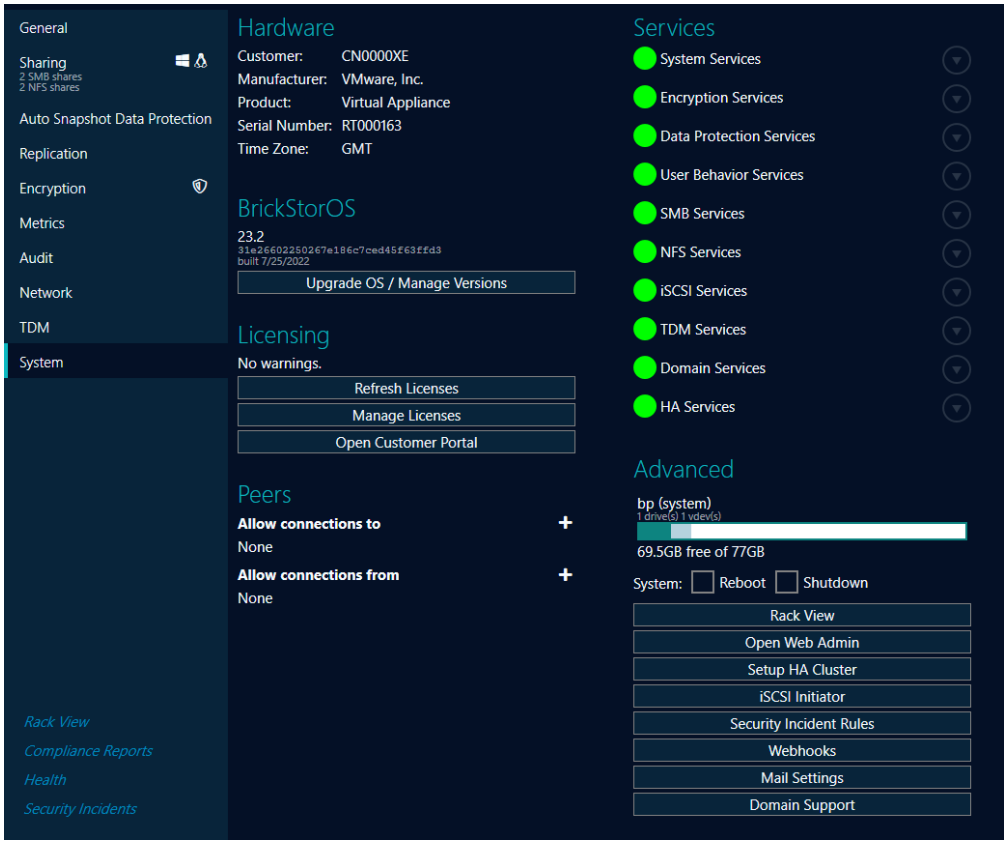

Installing the BrickStor SP on a Virtual Machine

1. The BrickStor SP Console will provide an **Admin URL**, as well as an **Access Code** located at the bottom of the home page.

2. With an open internet browser, enter the Admin URL into the url and press **Enter**.
3. A screen will prompt the input of an Access Code.

NOTE The Access Code is located within the BrickStor SP Console.

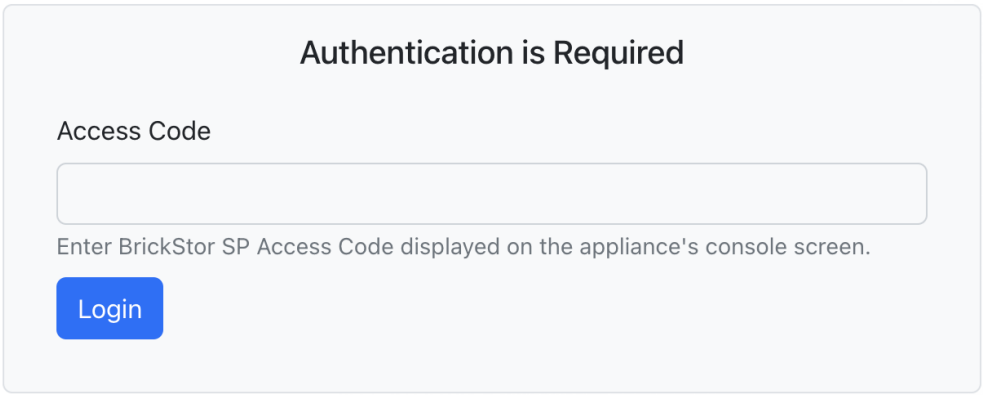

4. Once the Access Code has been entered, click the **Login** button.
5. The installer **Home** page will present.

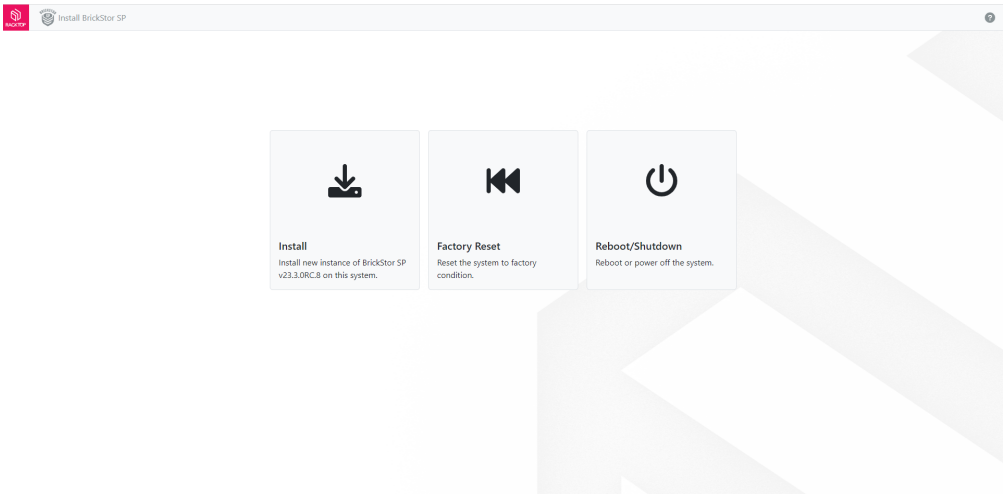

6. This page displays three buttons:
   - Install - Install a new instance of the BrickStor SP Manager.
   - Factory Reset - Reset the system to factory condition.
   - Reboot/Shutdown - Reboot or Power off the System.
7. The End User License Agreement screen will present.

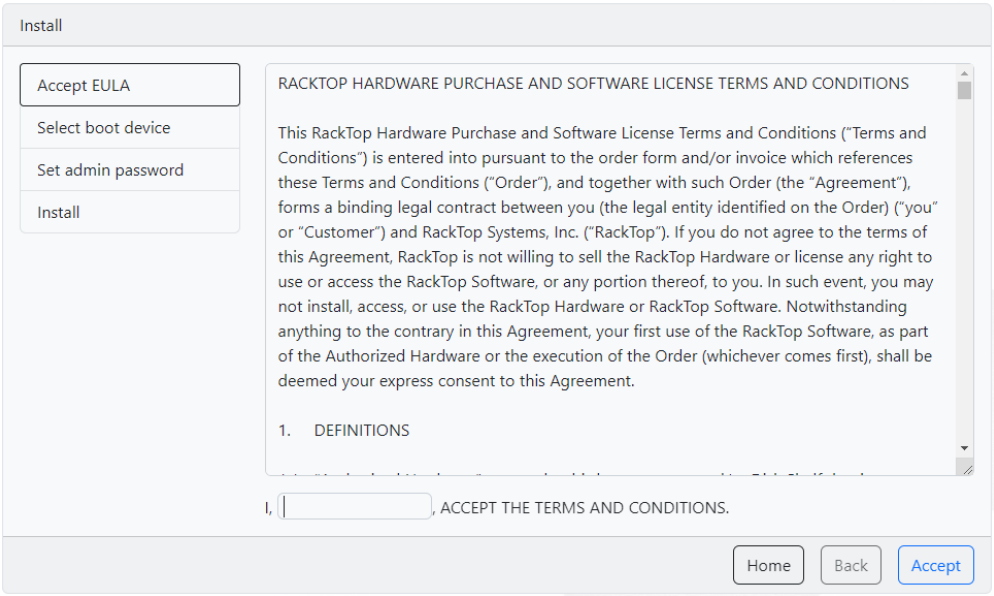

8. After reading the End User License Agreement, click the empty field provided below, and **enter the name** of the EULA acceptor.

```
I [FIRSTNAME LASTNAME] ACCEPT THE TERMS AND CONDITIONS.
```

9. Once finished, click the **Accept** button to proceed.
10. A green checkmark next to **Accept EULA** will be present upon success.
11. The **Select Boot Device** screen will present.

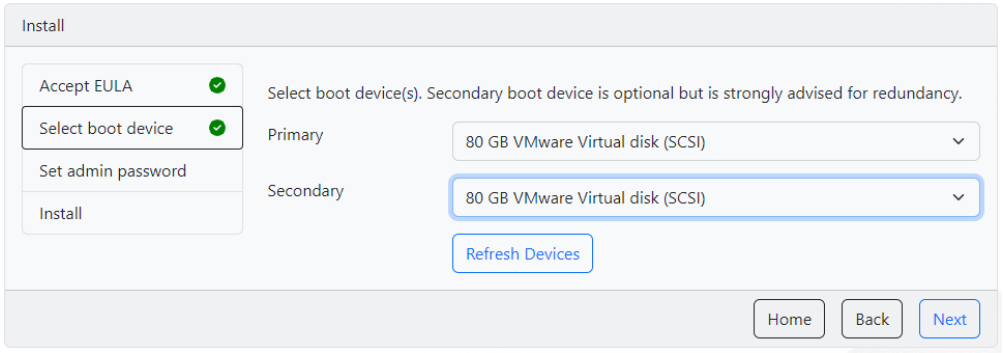

12. Select the Primary and Secondary boot devices by **clicking the provided fields** and selecting from the provided options.

NOTE	If the correct boot device option is not present, click the **Refresh Devices** button, and try again.

13. Once the correct primary and secondary boot devices have been selected, click the **Next** button.

NOTE	For disk selection only 1 disk is needed for Virtual Machine installations. For bare-metal installations, two drives are strongly encouraged. When secondary disk is selected, a drive mirror (RAID-1) will be created for redundancy.
CAUTION	Secondary drives must be of the same type and capacity. SAS, SATA and NVMe device types are supported, but cannot be mixed.

14. A green checkmark next to **Select boot device** will be present upon success.
15. The Set Admin Password screen will present.

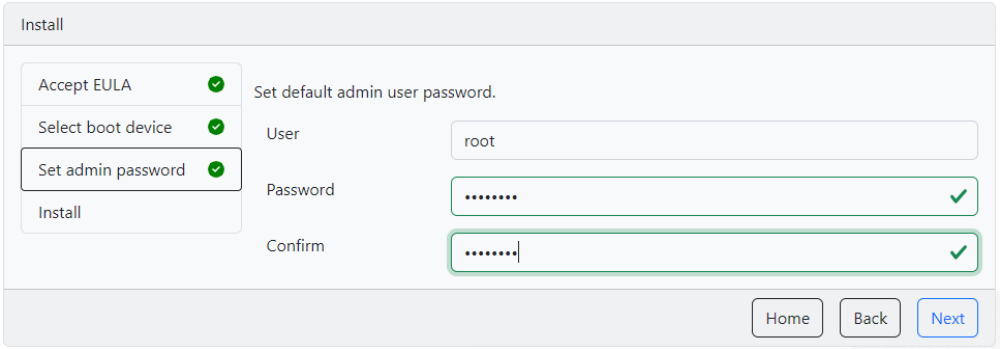

16. The user **root** will populate in the provided **User** field.
17. Enter the **Password** in the provided field.
18. Enter the **same** password in the following field to **Confirm**.
19. Once finished, click the **Next** button.
20. A green checkmark next to **Set admin password** will be present upon success.
21. The **Install** screen will present.

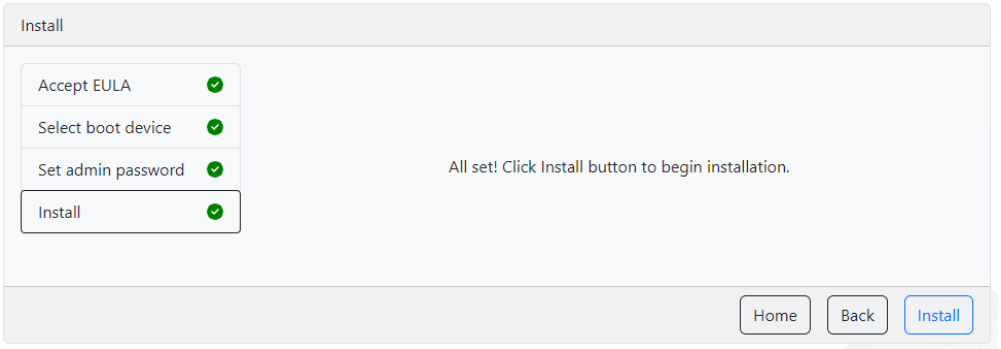

22. Click the **Install** button to begin installation.
23. A progress bar, and description of the installation process will present.

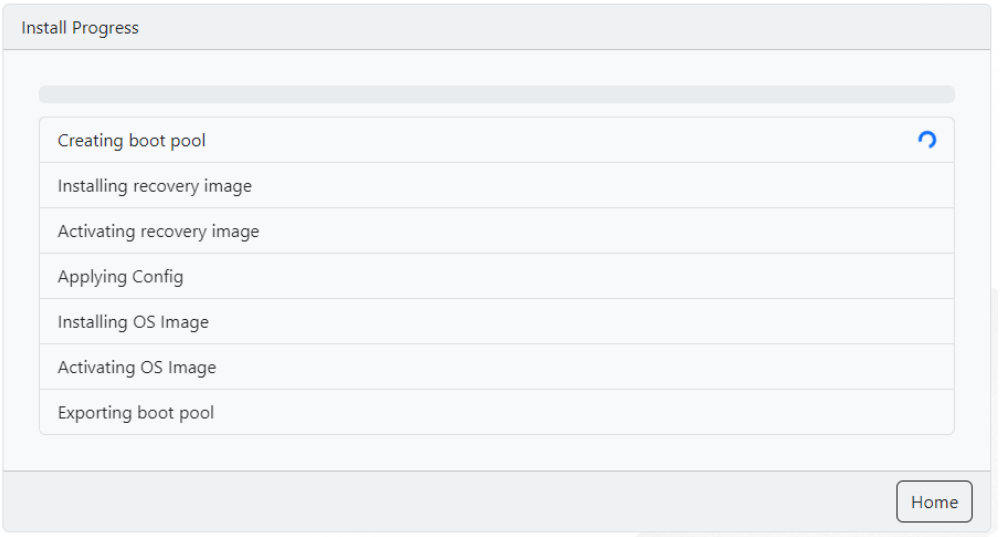

24. Once the progress bar has filled, a prompt denoting a system reboot will present.
25. A prompt will present denoting successful installation, click the **Reboot** button.

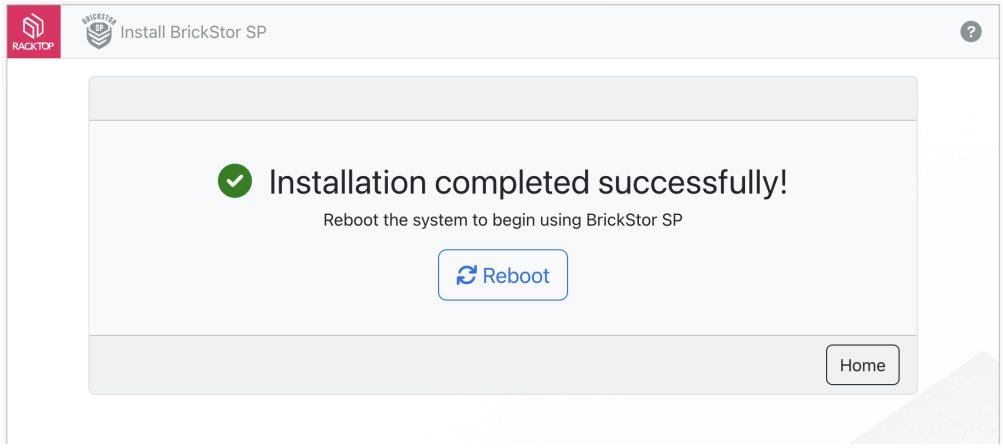

NOTE When Installing a new instance of the BrickStor SP Manager, click the **Install** button, and repeat the above installation steps.

NOTE To Reboot the System, click the **Reboot** button on the home scree, then click **Reboot**.

Joining Active Directory

BrickStor SP appliance is capable of integrating into an existing Active Directory environment, which allows for share permissions and administration delegation to reference users and groups in Active Directory. To associate user and group permissions on shares, directories and files with user and group objects in the Active Directory, the appliance **must** first be joined to the Active Directory Domain.

Active Directory Join Prerequisites

A number of basic requirements must be met before domain join can succeed.

- Working Domain Name Service (name resolution), For resiliency, having two or more Domain Controllers is strongly advised.
- Configured NTP. Accurate and reliable time-keeping with clocks synchronized between the BrickStor SP appliance and Domain Controllers.
- Endpoints/client connection to BSR using AD authentication for SMB/NFS services are also required to have clock synched with AD.
- Account username/password with proper access to perform the Active Directory joining.
- Fully Qualified Domain Name (FQDN) (ex: example.com) of the domain to be joined.

Configuring NTP

It is a best practice to configure domain joined BrickStor SP appliance with specific domain group(s) to allow management access. This will allow group members to access BrickStor SP Manager using domain logins.

Run `setup`. The main menu will present itself, select option 5, *Configure NTP settings*.

```
$ setup

RackTop Cyberconverged NAS
Setup Utility
Copyright 2021 RackTop Systems, Inc.

  Main Menu

  1. Configure RMM interface.
  2. Configure nodename.
  3. Configure network interface.
  4. Configure aggregate network interface.
  5. Configure NTP settings. <- Select this option
  6. Configure DNS settings.
  7. Disable system service connections to the Internet.
  8. Configure TimeZone.
  9. Restart appliance.
 10. System Information and Administration.
 11. Exit Setup Utility.

 Select menu option and press enter or press enter to exit.
 Use CTRL-C to exit at anytime.
```

Next, select option 1, *View current NTP settings*. This will present the currently configured NTP server(s). On a brand new system the output is going to resemble the following:

```
Missing NTP servers, consider adding at least one.

Press enter to continue.
```

If there are already configured NTP servers and the time is synchronized with the Active Directory Domain Controllers, the system is ready to join to the Active Directory. Otherwise, be sure to configure NTP servers.

WARNING	For VM deployments: Avoid using configuration to synchronize VM clock with the hypervisor and use NTP instead.

If changes are necessary:

- Press **Enter**
- Run `setup` again, the main menu will present itself.
- Select option 5, *Configure NTP settings*.
- Instead of option 1, after choosing to configure NTP settings, select option 2, *Configure NTP settings*.
- A prompt will direct the input of the IP address of an NTP server; alternatively, it is possible to use DNS names as well.

- A prompt will follow to confirm these inputs.

NOTE option 2, *Configure NTP settings* is additive, each time it is chosen and an address or DNS name is entered, this address or DNS name will be appended to the list. This list may or may not be empty, It is possible to adjust this list with option 3, *Remove NTP Server*.

```
RackTop Cyberconverged NAS
Setup Utility
Copyright 2021 RackTop Systems, Inc.

  NTP Configuration Menu

  1. View current NTP settings.
  2. Configure NTP settings.
  3. Remove NTP Server.
  4. Verify NTP Server request and synchronization.

  Please select menu option and press enter or press enter to return to main menu.
  2

  NTP Server IP Address: ad1.example.com

  Is the above NTP Server IP Address correct? (options: y or n):
```

NOTE Microsoft Active Directory default maximum tolerance for computer clock synchronization is 5 minutes, however, it is a configurable setting so it can vary in some environments.

NOTE In many cases Active Directory Domain Controllers are also DNS and NTP servers.

The following steps will outline the process of joining the Active Directory using the BrickStor SP Manager:

- With the BrickStor SP Manager open, navigate to the **System** tab.

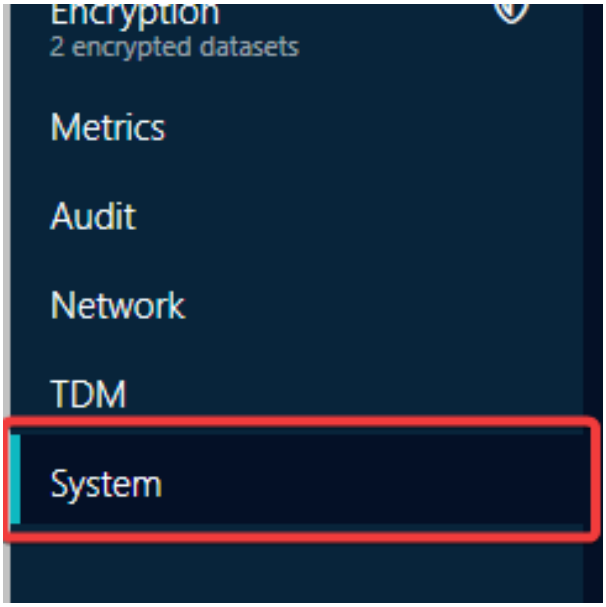

- Navigate to the **Domain Support** Tab.

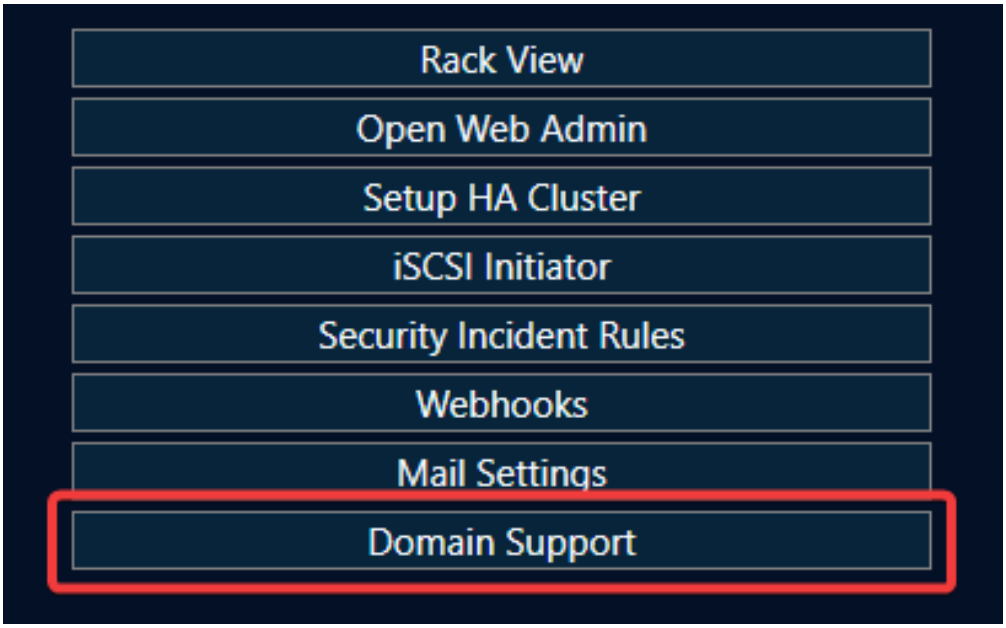

- Check the box to **Modify/Repair Domain Membership**.
- Enter the intended domain name to be joined.
- Enter the user name of the domain user with sufficient domain join privileges.

Managing BrickStor SP with AD Users

BrickStor SP can be configured to allow specified Active Directory groups to login to BrickStor SP Manager and administer the appliance. This is useful in order to track changes made to a BrickStor SP via the [Audit Log].

- Ensure prerequisites are satisfied concerning the Active Directory Server.
  - Join system to the domain, or, ensure the system has already been joined to the domain.

The following steps will outline the process of joining the Active Directory via the BrickStor SP Manager:

- With the BrickStor SP Manager open, navigate to the **System** tab.

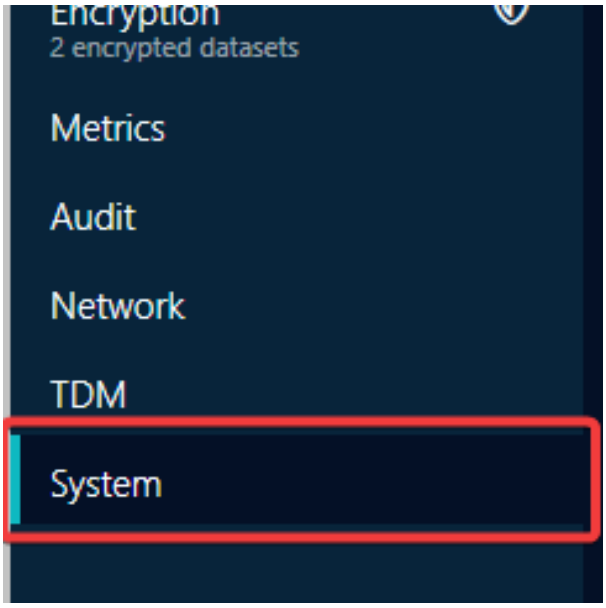

- Navigate to the **Domain Support** Tab.

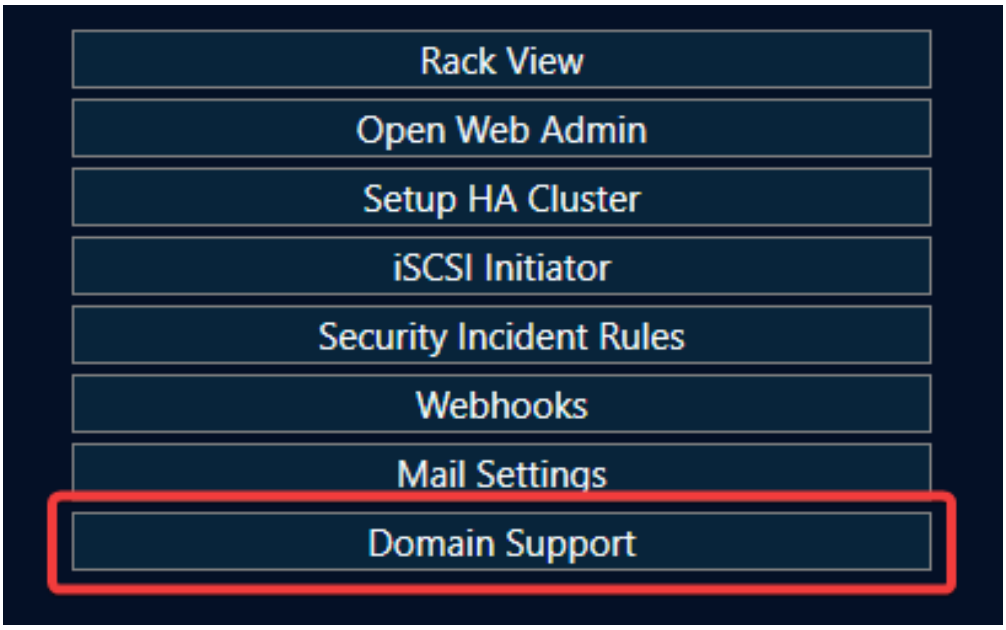

- Click to expand the **Admin Group** section.

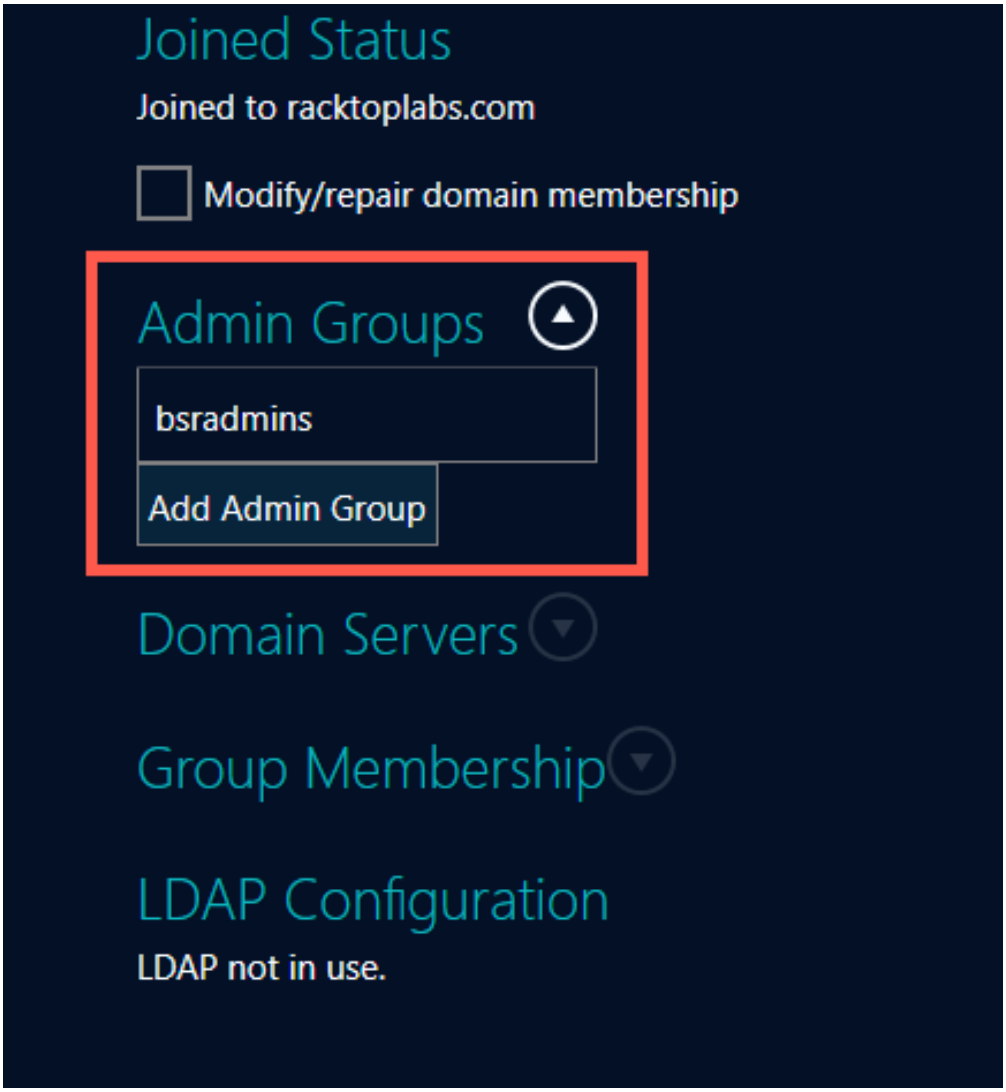

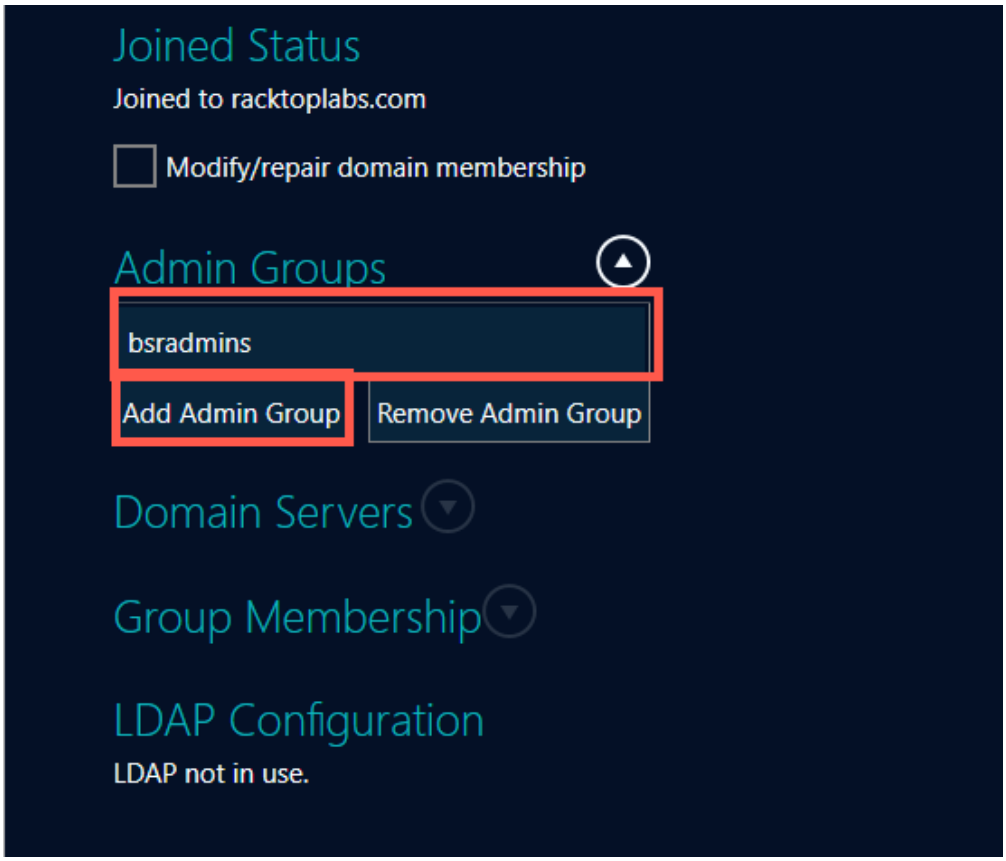

- Click **Add Admin Group**
- Search for and select the group that is to be granted Active Directory Credentials.
- Click **Add Admin Group**

NOTE

A group may also be removed from Active Directory permissions. To do this, follow the above steps to select the Admin Group that is to be removed. Once selected, click **Remove Admin Group**.

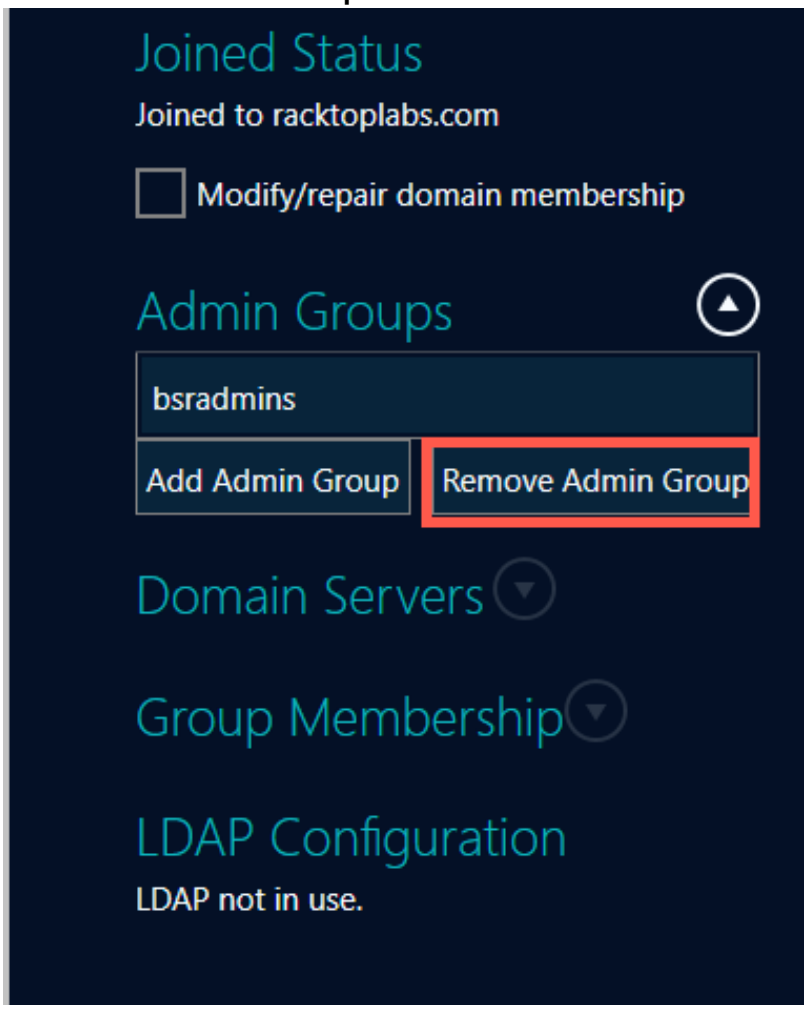

Data Protection

Data Protection Information

Data protection encompasses point-in-time snapshots of datasets and replication of these snapshots. 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 Data Protection Information

To view BrickStor SP 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 **Auto Snapshot Data Protection** tab.

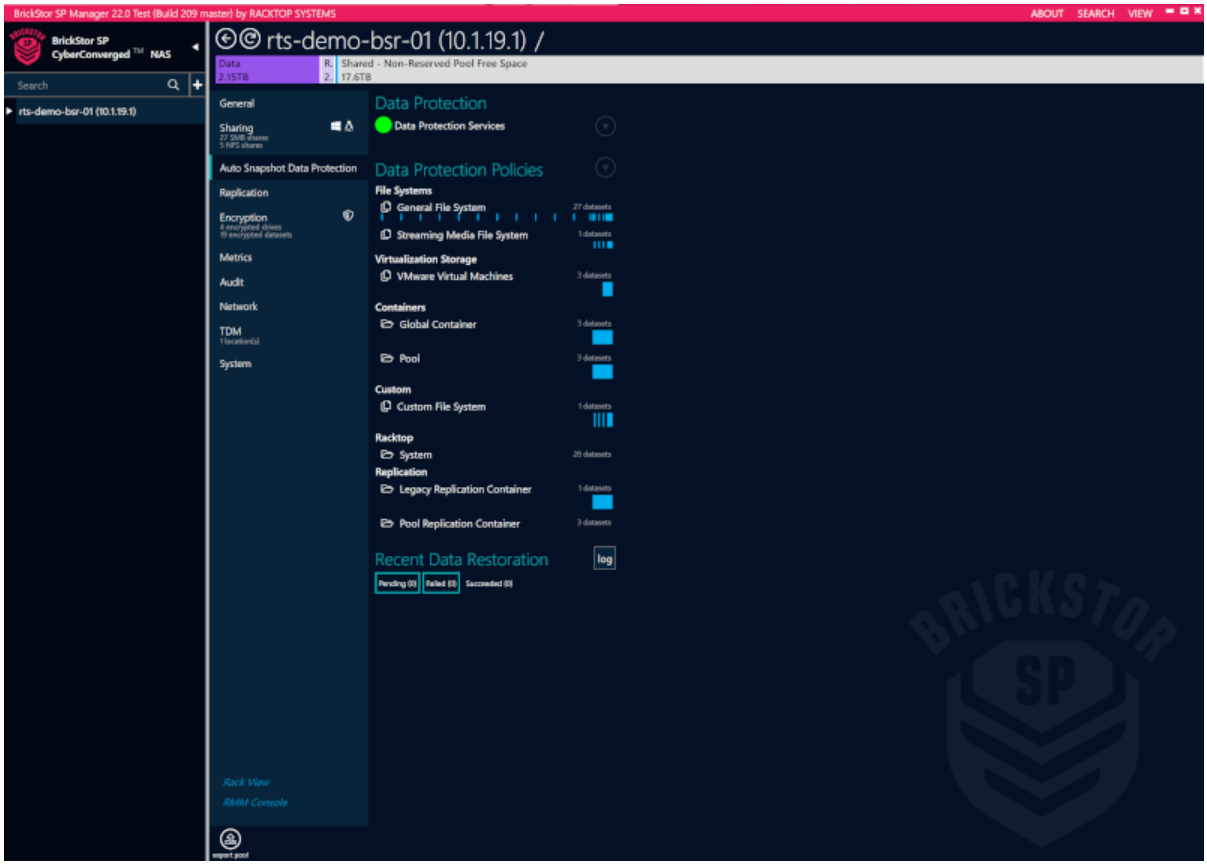

Additionally, one can also select a specific dataset in the navigation pane and select the **Auto Snapshot Data Protection** tab to view the policies for a specific dataset. Each dataset can either

have their own individual policy, or the dataset can inherit the policy from its parent dataset.

Configuring Data Protection

Select the dataset to modify and navigate to the **Auto Snapshot Data Protection** page as described above. The following elements on the Auto Snapshot Data Protection page can be used to configure the data protection policy:

- The **log** button takes you to a screen which provides detailed logs of snapshot activity.
- The next run, last run, and last status displays indicate the corresponding values for the selected dataset.
- The protection policy dropdown allows the choice of whether to customize the policy, or to use the default policy for the storage profile.
- Next to the protection policy dropdown is a small button for indicating and toggling inheritance. This defines whether the snapshot policy will be inherited from the parent dataset or be independent.
- The On/Off switch

is used to enable or disable snapshots completely on the dataset.
- The bar is a graphical representation of the number of snapshots you will have with the policy over the course of a year.
- Snapshots can be configured to occur at a particular frequency, and each frequency has a customizable retention period.

In this example:

The first line indicates that snapshots will be taken every 15 minutes and they will be retained for one day. This drop-down can be used to set the frequency within a 24 hour period. Changing the retention will preserve storage space by deleting snapshots that are older than the setting. For example, leaving the frequency at 15 minutes and setting the retention value to 1 day, there will always be snapshots available going back 24 hours but any snapshot older than that will be removed. The retention setting can be set up to 99 days. Beyond 99 days, the value is infinite—meaning snapshots will never expire.

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.

NOTE This walkthrough primarily covers hardware-centric deployments, and may not represent effectivity with virtual deployments.

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 know 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 NFS or SMB. 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.

NOTE This walkthrough primarily covers hardware-centric deployments, and may not represent effectivity with virtual deployments.

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 ratio 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 1. 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)

Type	Description
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.

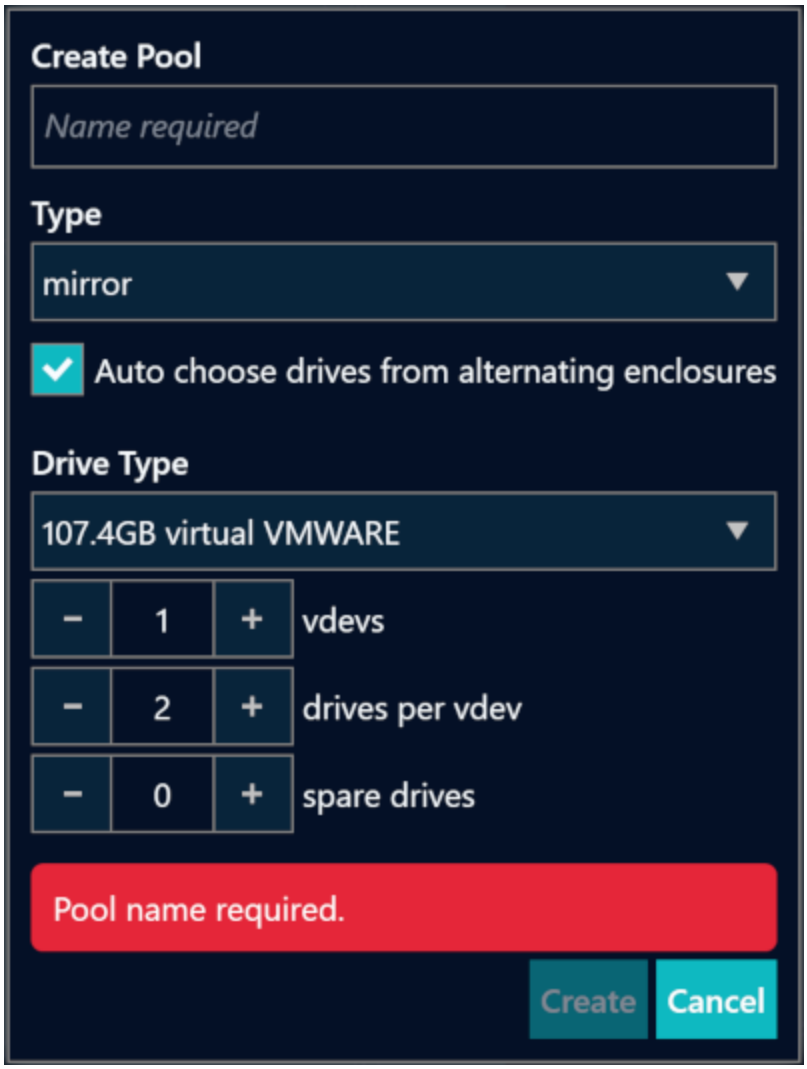

3. In the Create Pool dialog box, type a name for the pool.
4. Under **Type**, choose one of the following options:
   - mirror
   - raidz3
   - raidz2
   - raidz1
   - disk
5. Optionally, select to **Auto choose drives from alternating enclosures** if you want BrickStor SP Manager to select the drives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

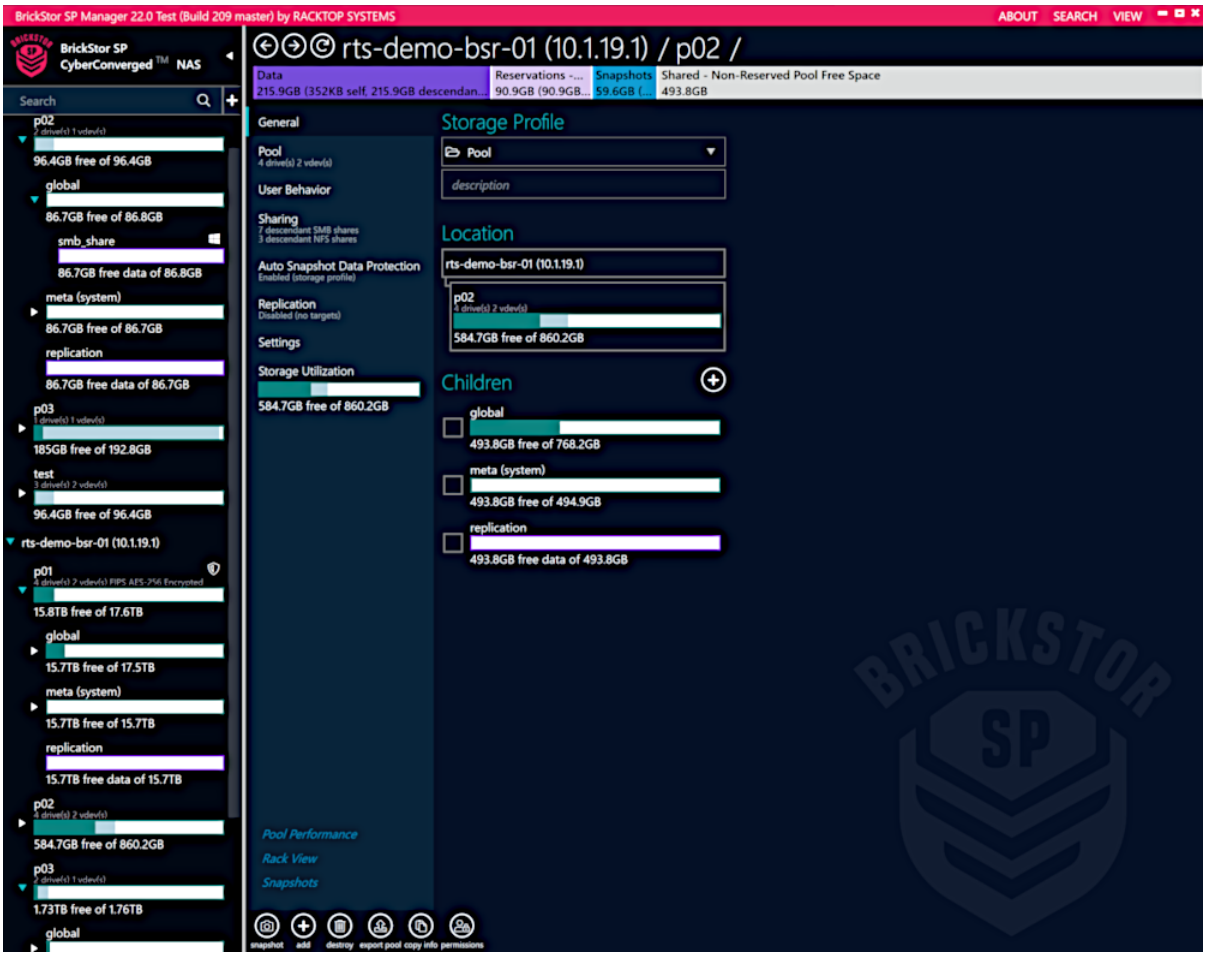

Managing Pools

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.

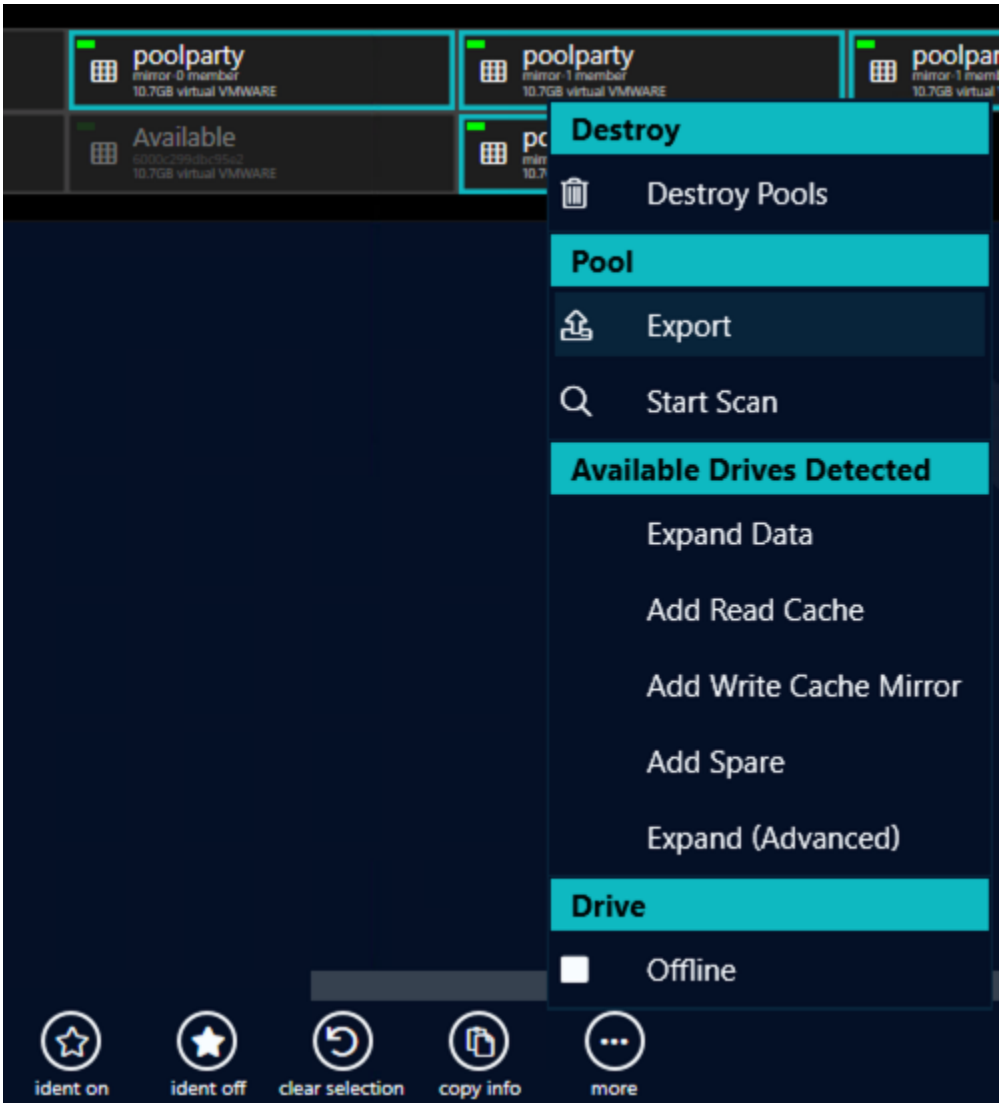

The second option is to select the pool from the Connections pane on the left-hand side of BrickStor SP Manager and click either the Expand Data, Add Read Cache, Add Write Cache, or Add Spare button under the Pool heading, depending on what you would like to add to expand the pool (will only appear if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

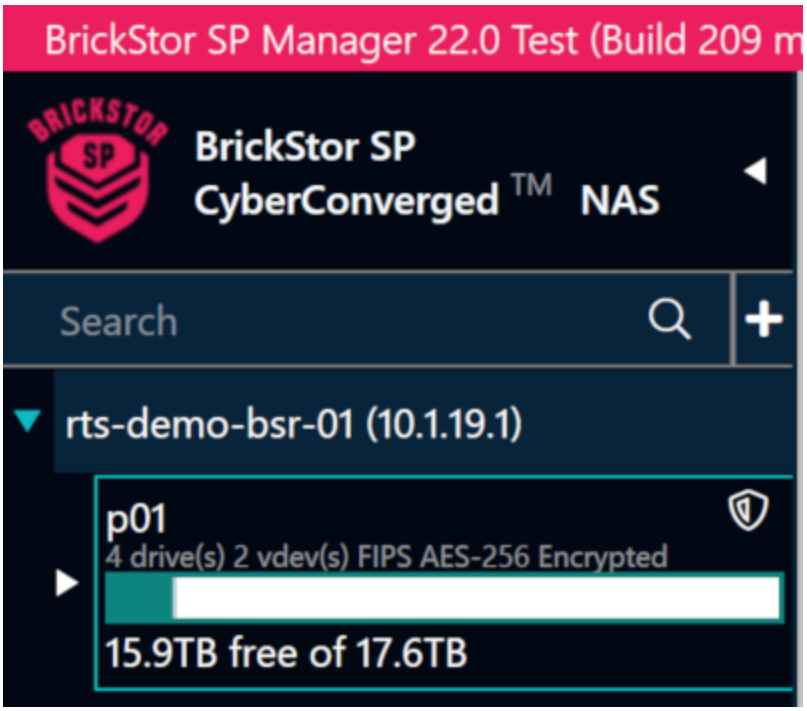

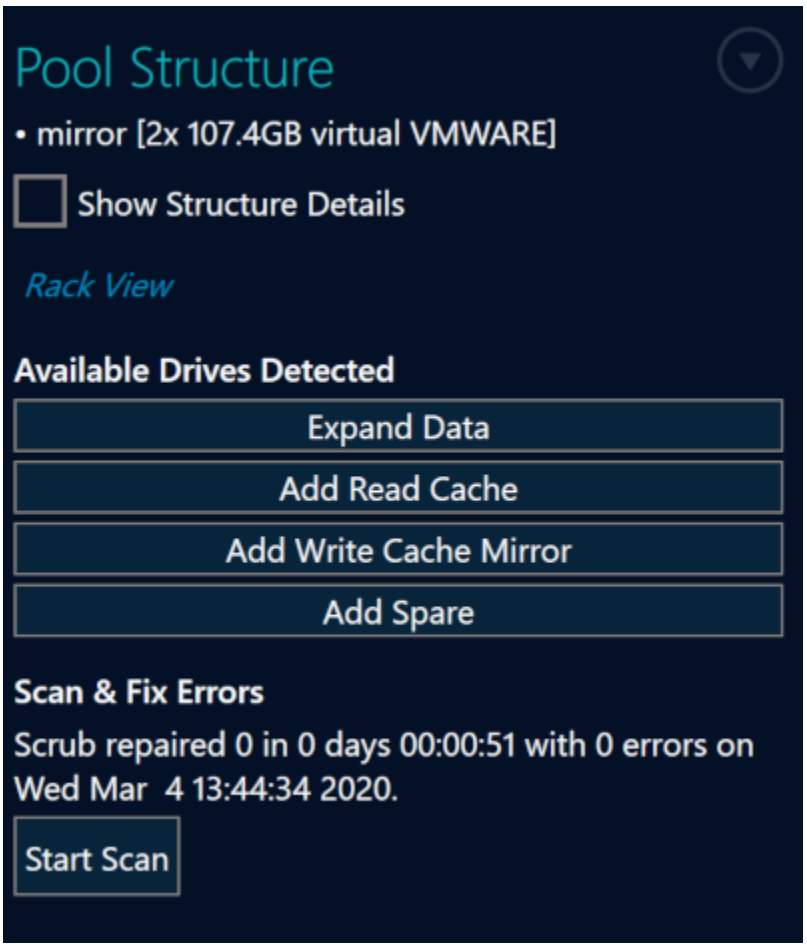

This will bring up the Expand Pool dialog box where you can choose to expand the pool by adding more vdevs, read and write caches, or spares. When the desired settings have been configured, click create to queue the change.

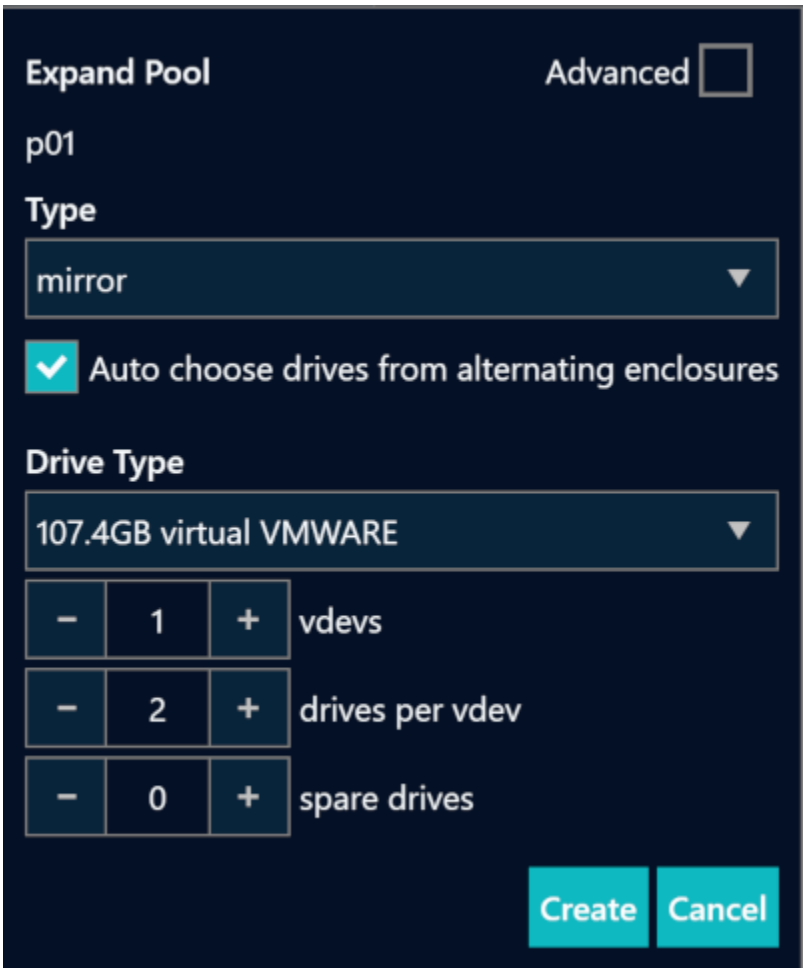

All changes in the queue will be indicated in Rack View and must be committed using the changes tab on the right side of BrickStor SP Manager.

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.

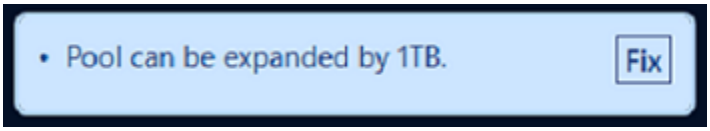

Replacing a Drive

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’.

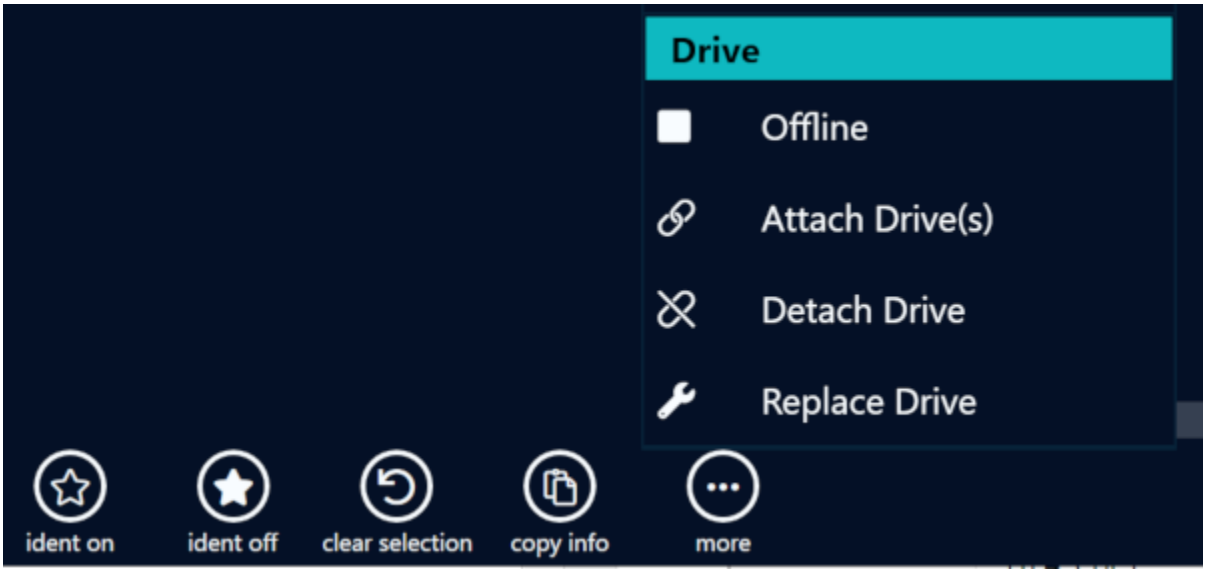

Or, if the drive is offline, you can navigate to the degraded pool in the Connections Pane on the left-hand side of the screen and click the Replace Drive button under the ‘Notable Vdevs & Drives’ heading.

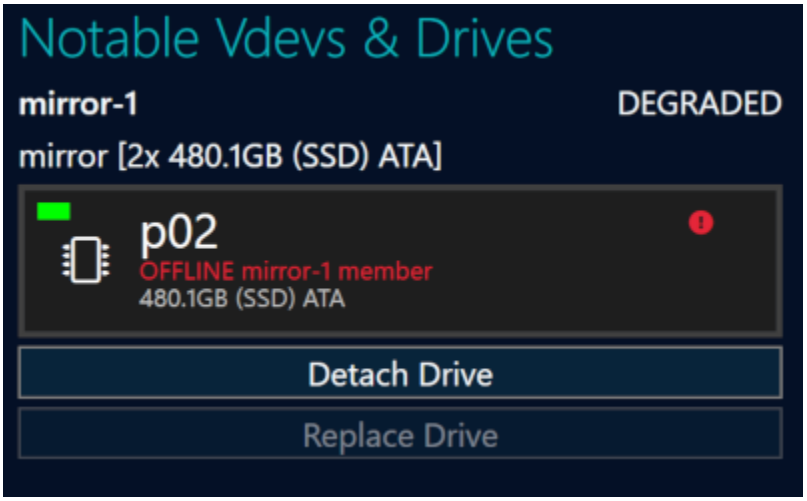

Selecting an offline drive from Rack View will also bring up actions that can be performed on it.

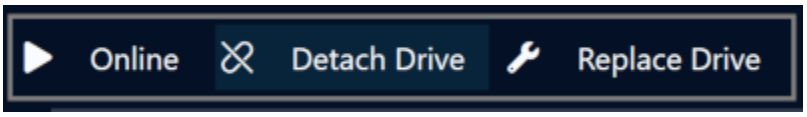

This will bring up the Replace Drive dialog box where you can select the drive to use as the replacement then click the Replace button to queue the change.

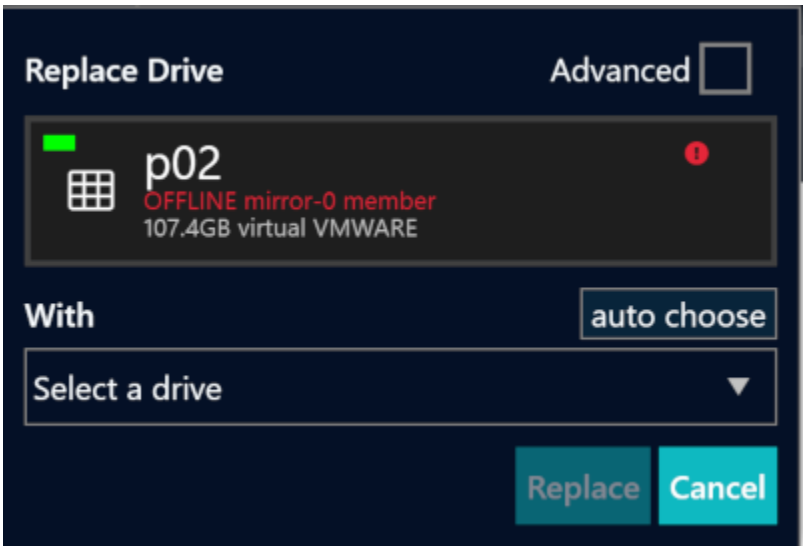

The change will be indicated in Rack View and will not be committed until the Commit Changes button is clicked on the Changes tab.

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.

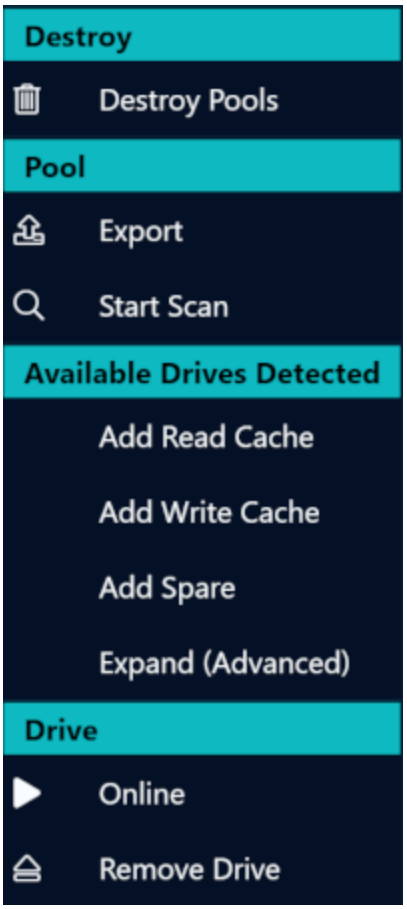

The change will be indicated in Rack View and will not be committed until you click the Commit Changes button in the Changes tab on the left-hand side.

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).

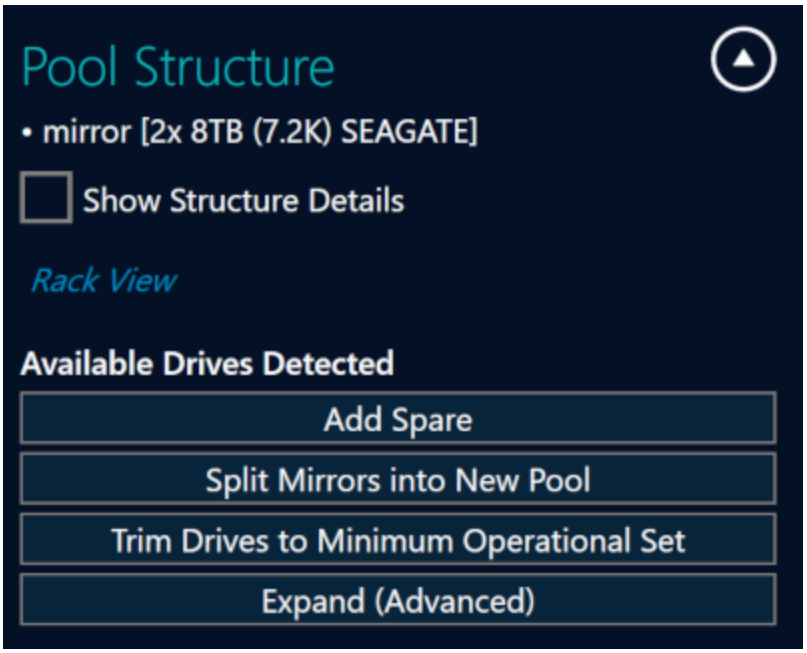

From the changes tab on the right-hand side you can change the name of the new pool that will result from the split and commit the changes with the Commit Changes button (by default the new pool created this way will be exported).

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.

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.

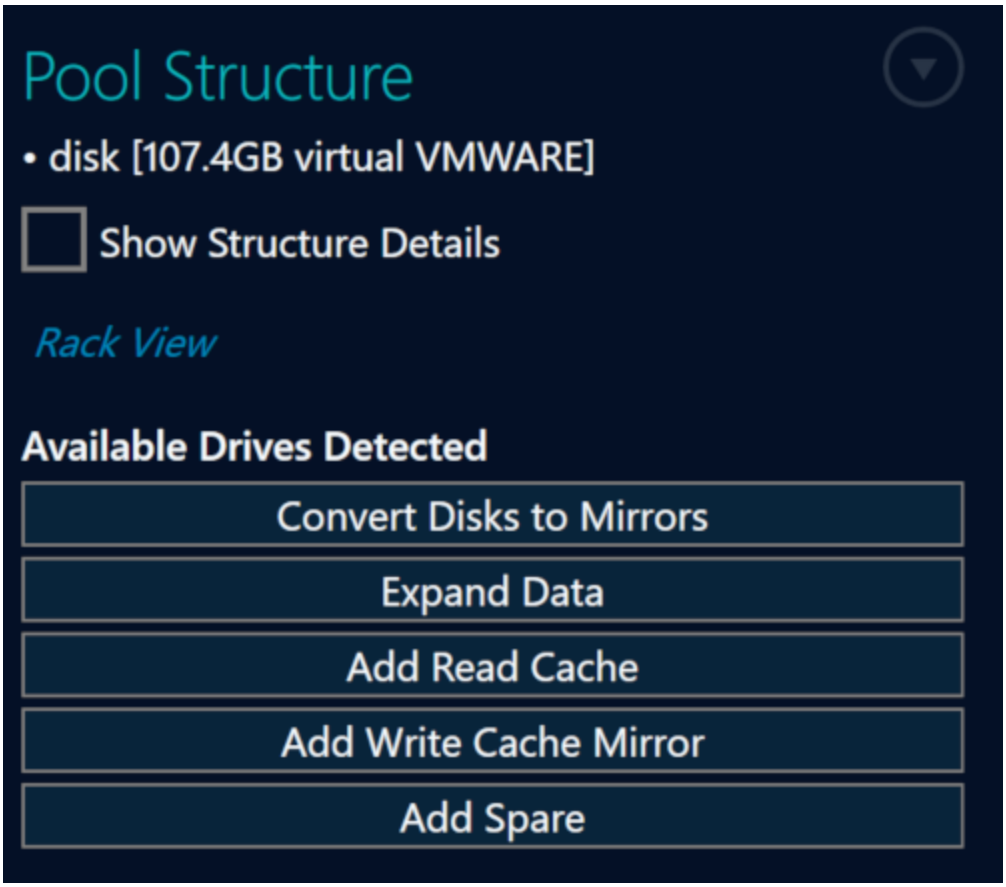

If done through Rack View, you will need to select the drive to attach yourself. When done through the pool's page it will select a drive for you automatically. The change will be indicated in Rack View and will not be committed until you click the Commit Changes button in the Changes tab on the right-hand 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
2
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.

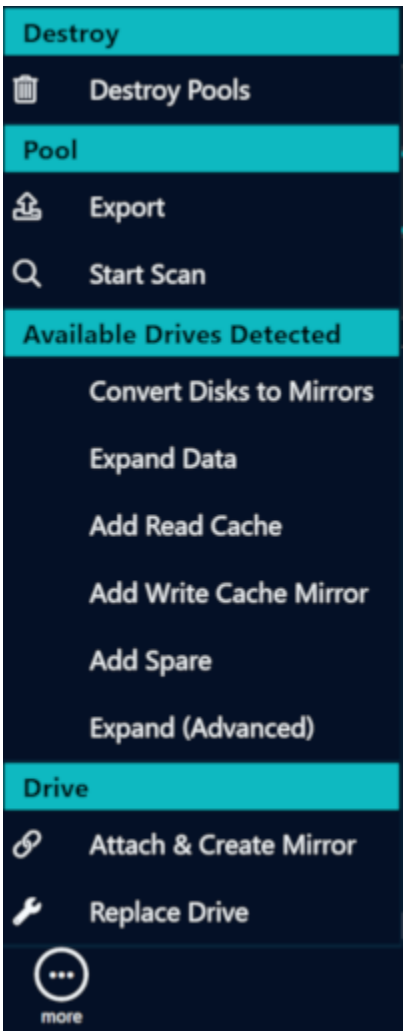

The button is also available on the Pool Tab.

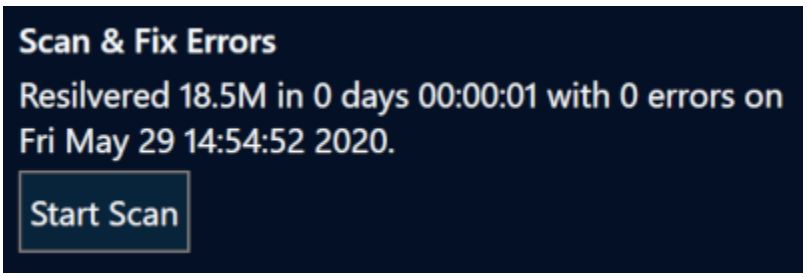

The scan will not be started until you click the Commit Changes button in the Changes tab on the left-hand side.

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.

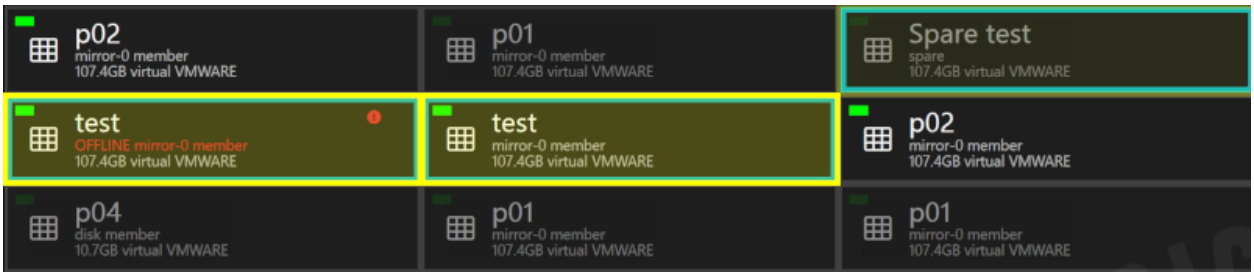

From the pool's screen on the Connections pane, the faulted drive will appear under Notable Vdevs & Drives. You can choose to promote the spare drive and detach the faulted drive from the pool, replace the faulted drive with another available drive on the system and return the spare to be a spare for 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.

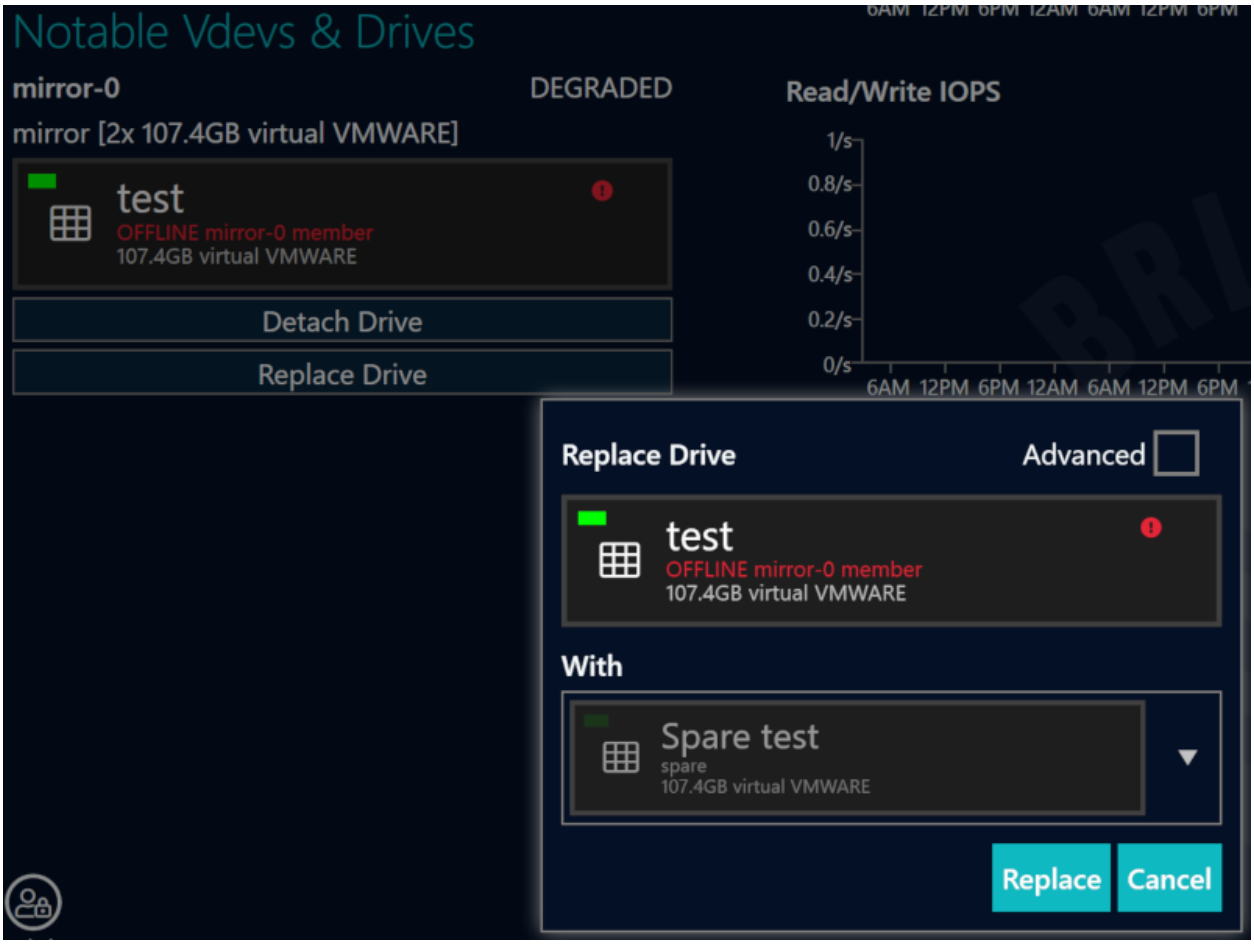

Each of these changes will require you to click the Commit Changes button in the Changes tab on the left-hand side to complete the action.

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.

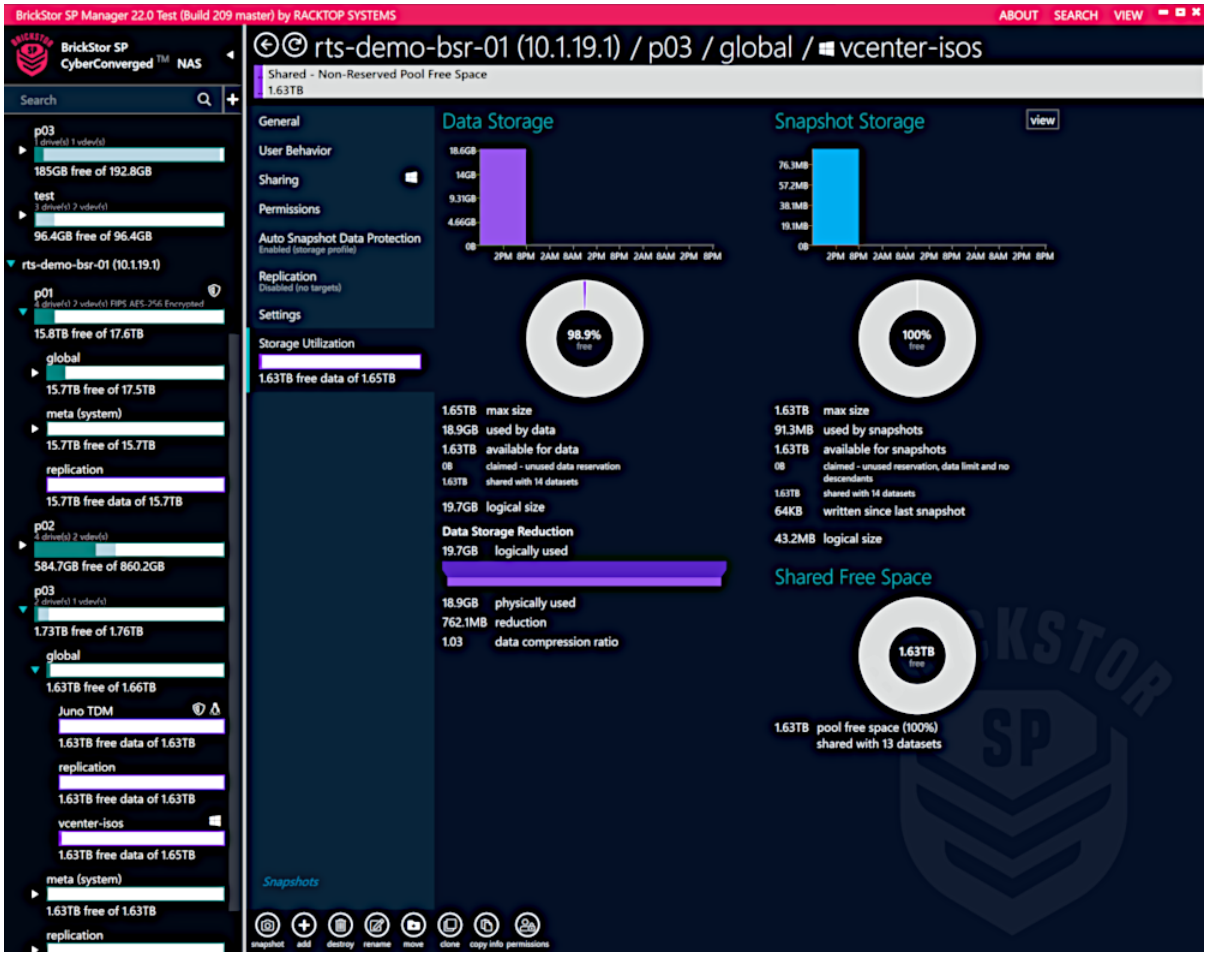

Pool Performance

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.

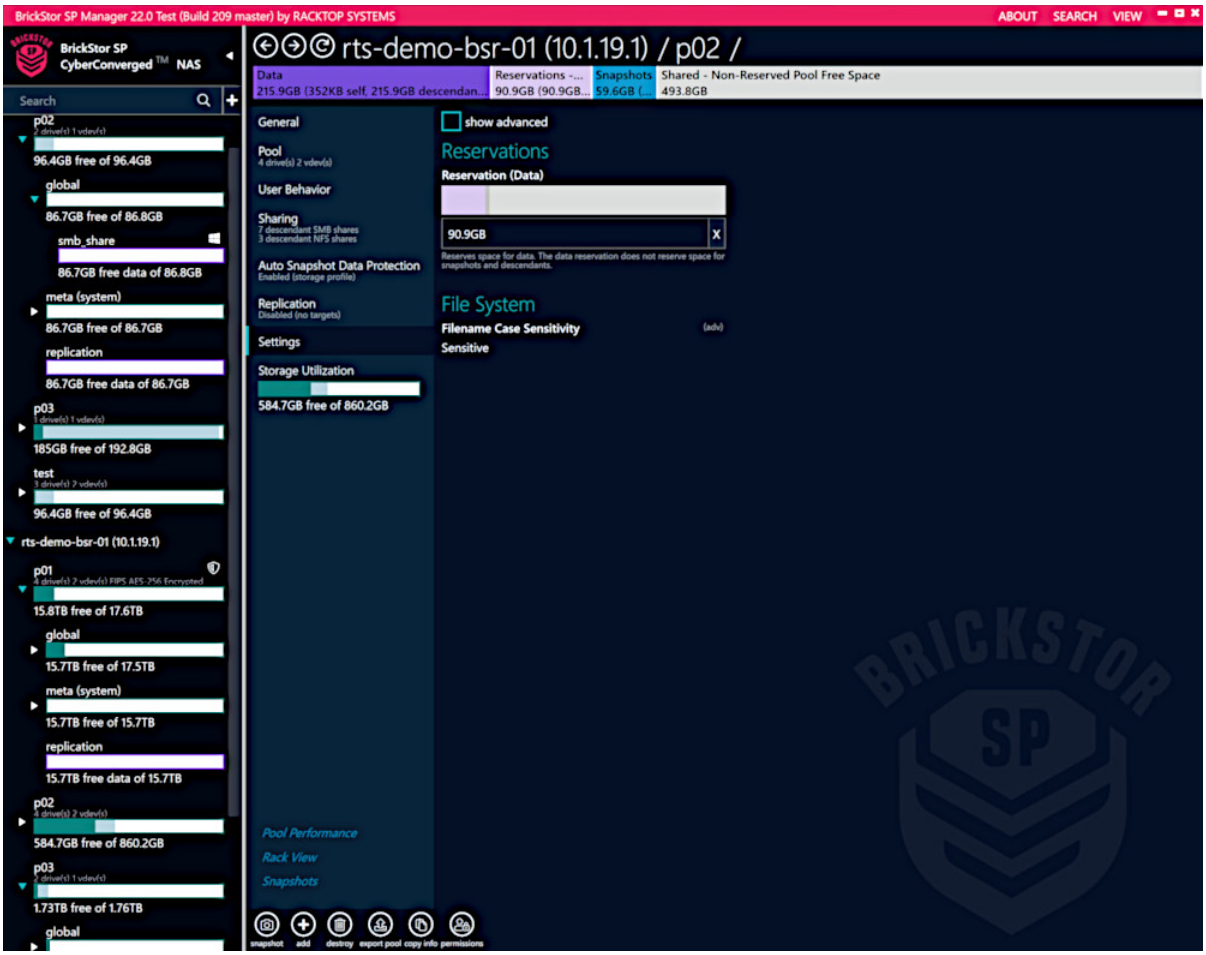

Destroying Pools

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.

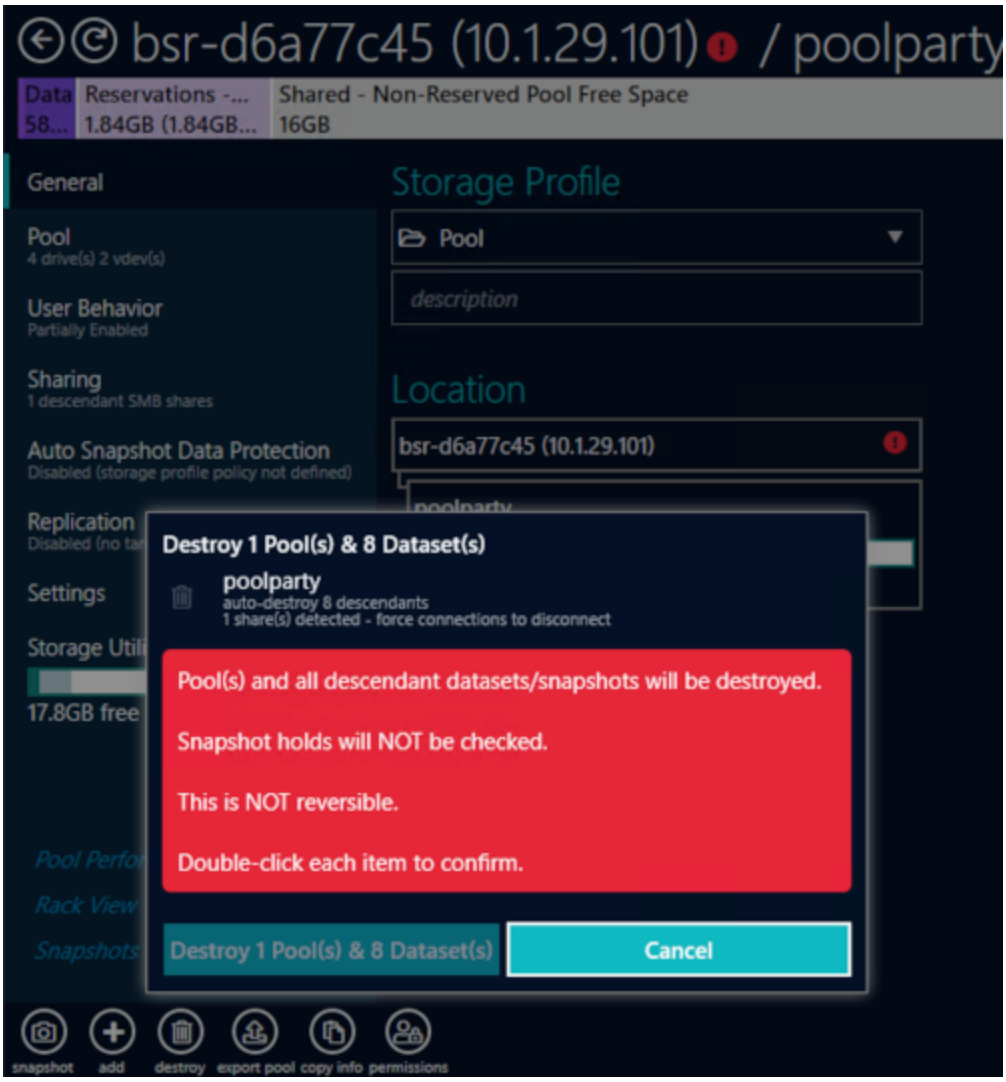

NOTE

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, the share name for the dataset.

Share Types

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

- SMB
- 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.

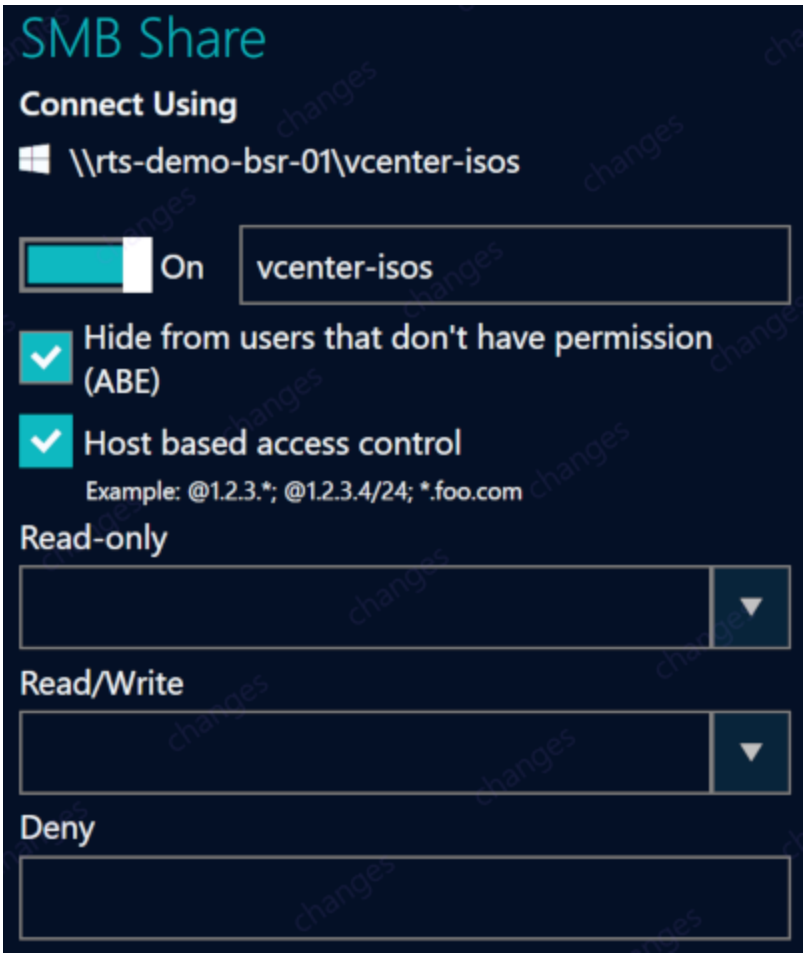

NFS

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.

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 a storage profile, based on your proposed workload.

   A storage profile defines a number of settings optimized for a particular kind of workload. Additionally, different storage profiles may have different settings available that are appropriate for that particular workload. This includes which methods are available to share a volume. Volume profiles (e.g. **General Volume**) create iSCSI volumes, while the profiles that do not contain 'Volume' create datasets that may be accessed using NFS and/or SMB (depending on the particular profile). For example, the **VMware Virtual Machines** storage profile can only be shared via NFS.

   Each storage profile also has an associated auto snapshot profile. The associated snapshot profile is the default snapshot policy for any datasets or volumes that are assigned the given dataset profile see Auto Snapshot Data Protection for more information.

   The available storage profiles are:

   - If you are setting up a File System:
     - General File System
     - Rendering
     - Streaming Media File System
     - 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
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

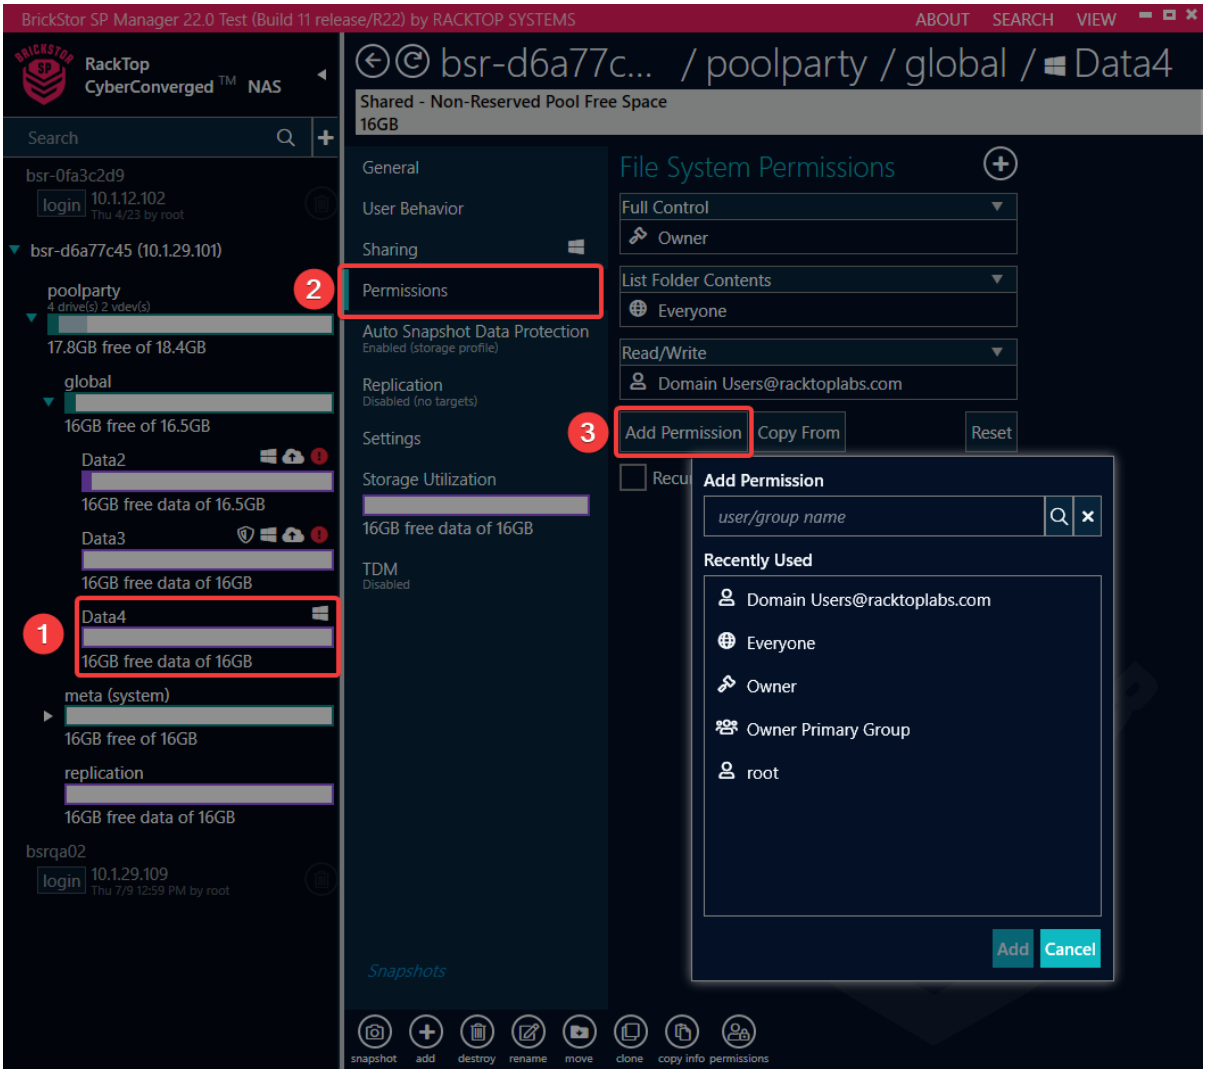

Using the Add Permission dialog, you can select previously used users or groups, or search for a user or group.

Add permissions search results

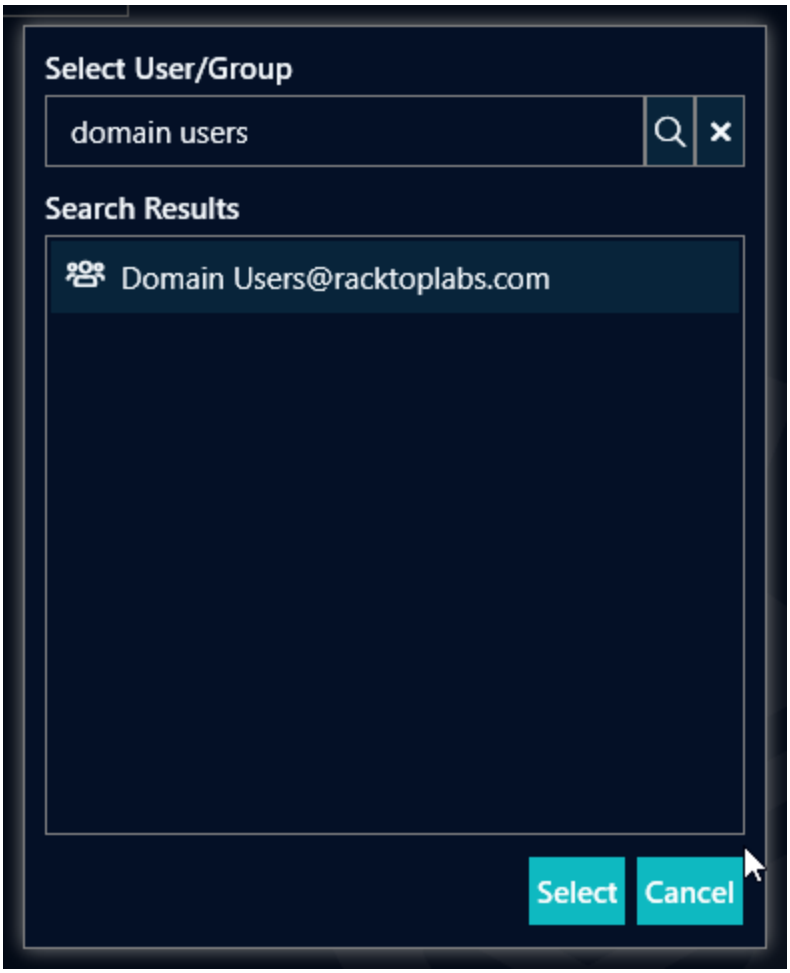

In the drop-down above the user or group, you can modify the type of permission. The default is Read/Write.

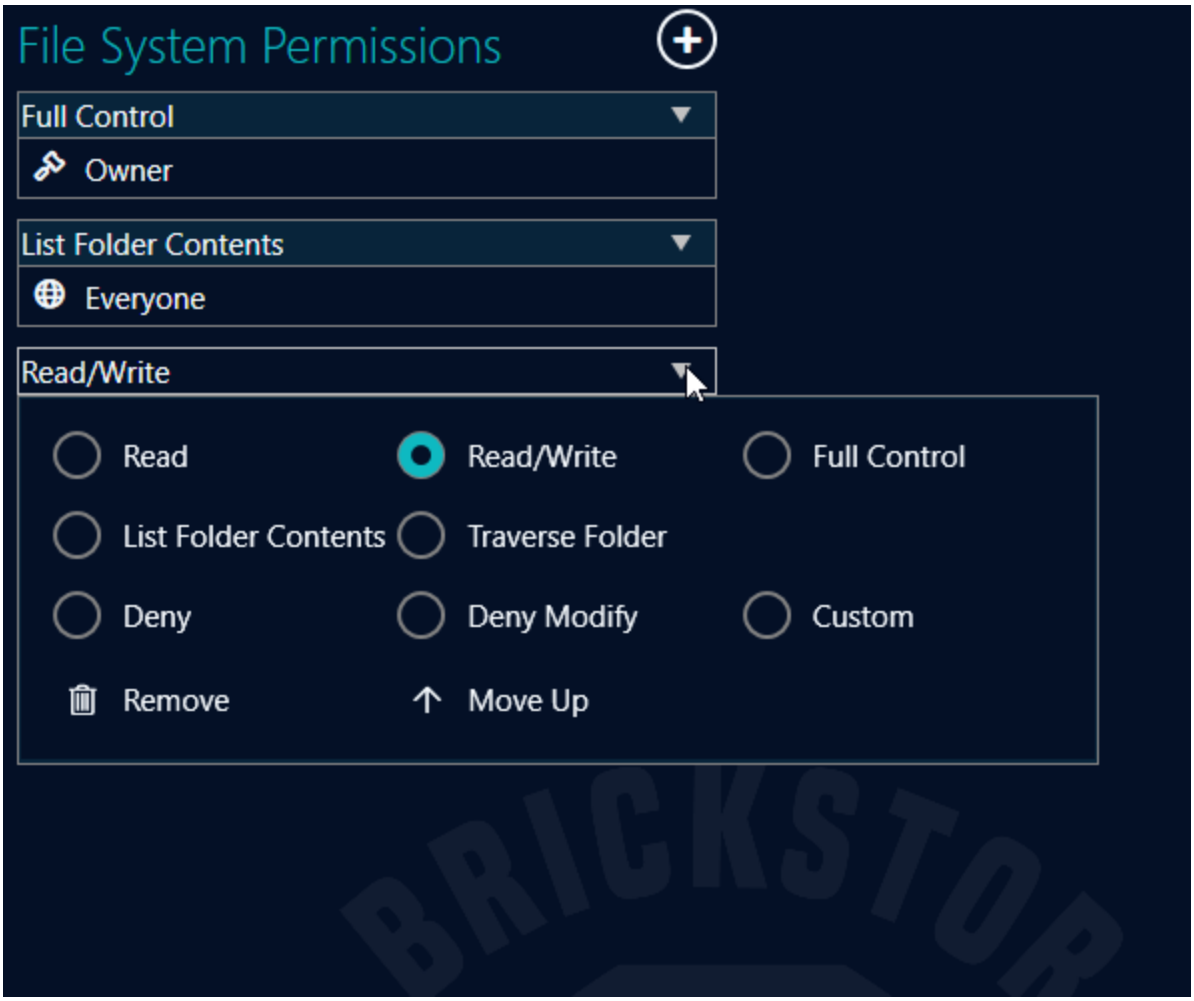

Additional options include recursively applying permissions or setting the new user or group as the owner. Once those choices are made, click the Commit button in the Changes pane to apply.

Choose permissions options and commit or undo

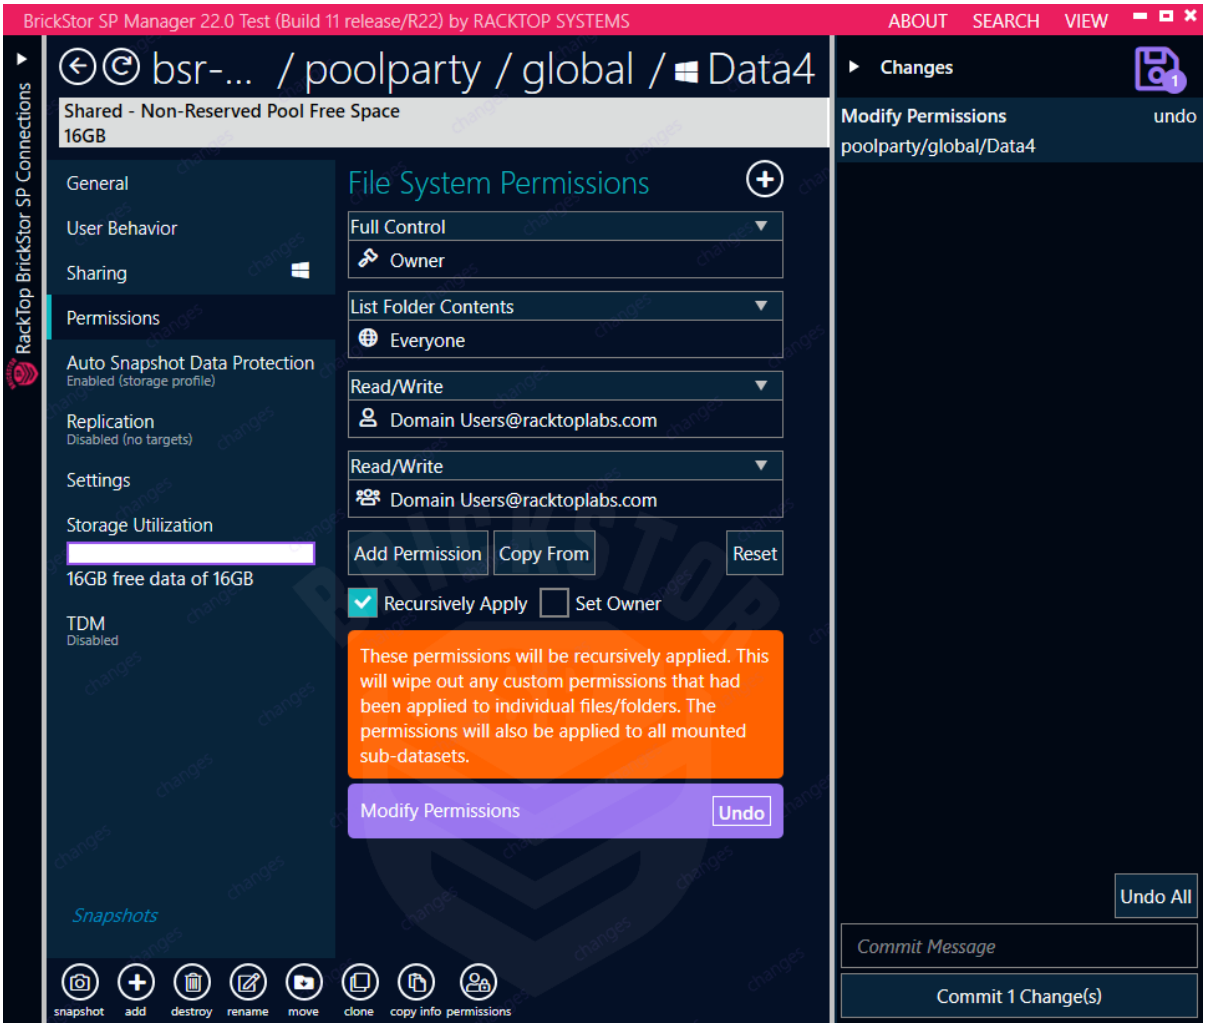

NOTE If the Recursively Apply box is not checked, permissions will only apply to newly created files and folders. Files created in existing folders will not be updated.

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 Bars

Throughout the Brickstor SP Manager, dataset bars are used to provide a color-coded quick view of the utilization of a dataset. The fraction of the bar that is filled in represents the amount of space being utilized. Since there are different types of utilization, different colors are used to indicate which category of utilization is shown.

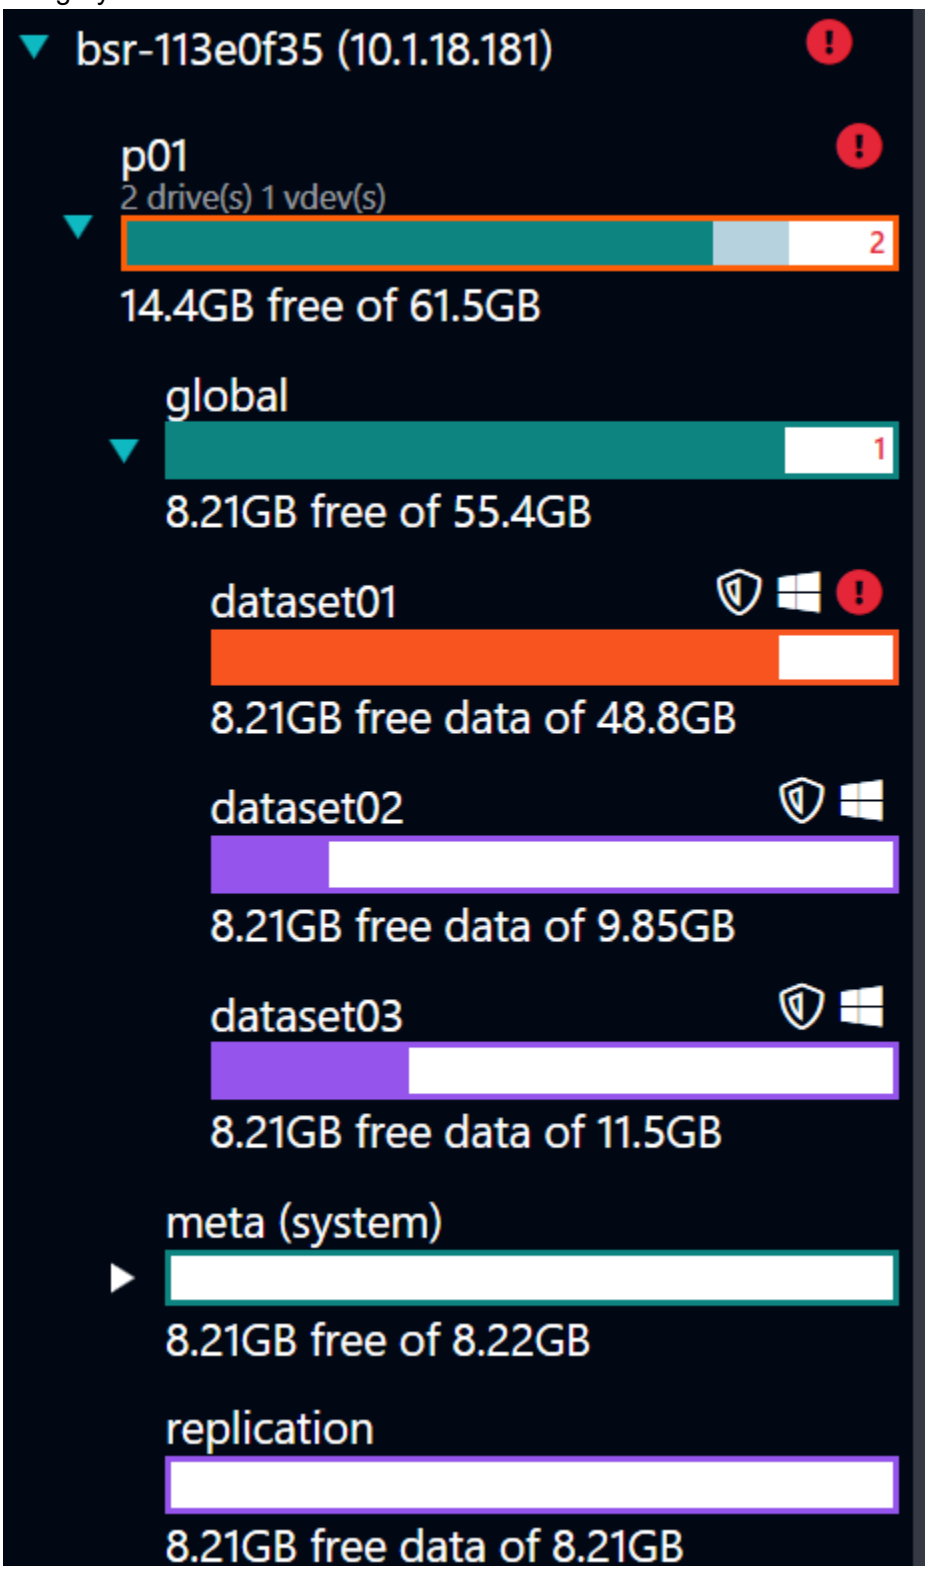

There are currently three categories of utilization using the following colors:

- Purple.

  The purple bar displays the ability to store data. The purple data bar is displayed if a data reservation has been set, a data quota has been set, the dataset has no children, or there is 25% or less free space for data (5% or less for archive storage profiles).

- Teal.

  The teal bar displays the ability to provision sub datasets. This is displayed if the dataset has children.

- Red/Orange

  The orange bar is displayed when the data set is low on storage.

Depending on quotas, refquotas, reservations and refreservations you could have different free space for each. Instead of showing two bars for each dataset, the UI attempts to show the relevant ones based on each datasets configuration and status. For example, container datasets generally show the teal color because they do not directly store data or snapshots. If the sub datasets do not have any children, the sub dataset will have a purple bar.

Dataset Storage Utilization

Storage Utilization allows you to view detailed 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.

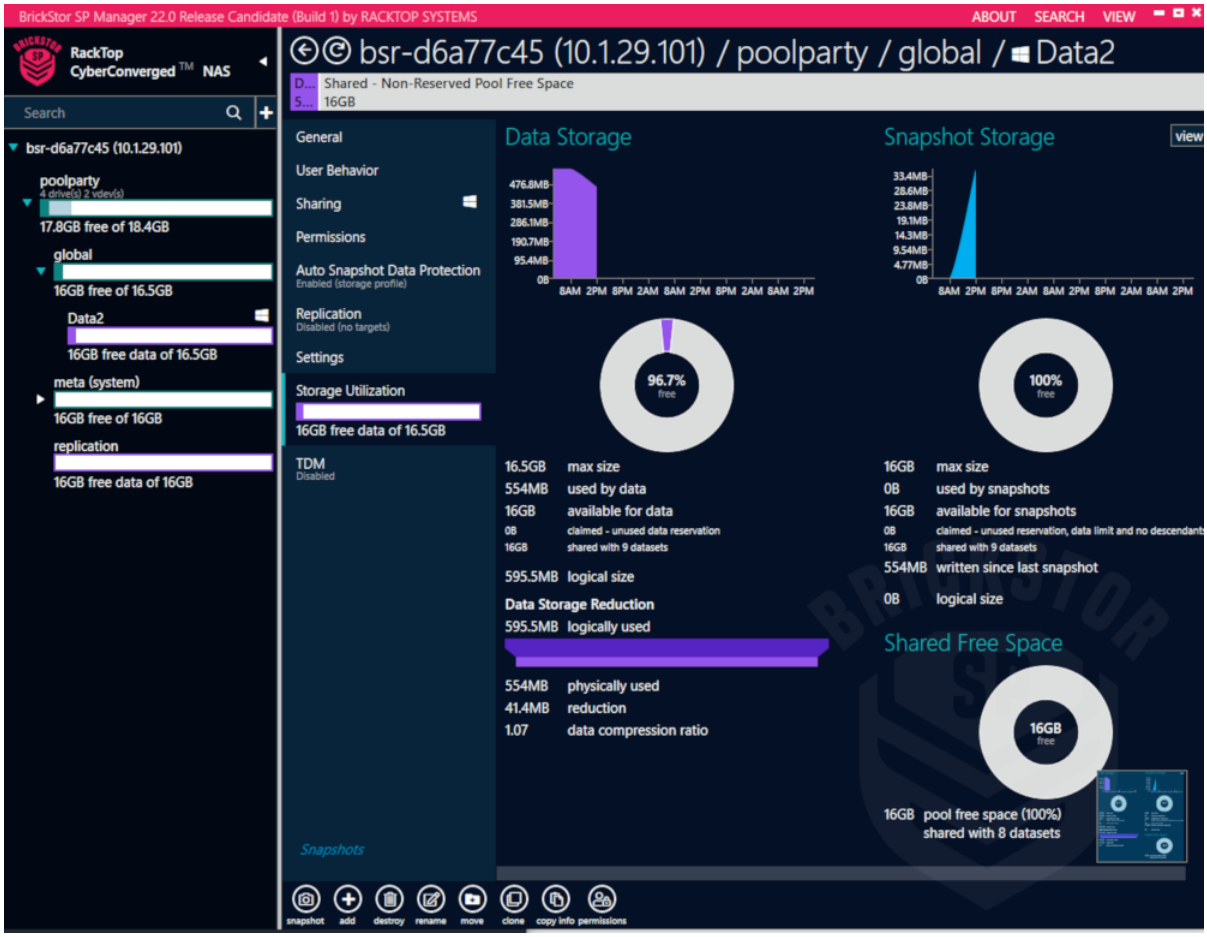

iSCSI

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.

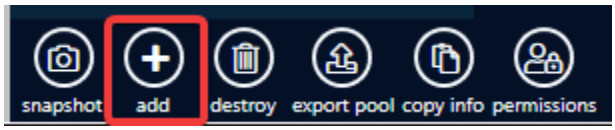

8. In the Create Dataset dialog box, type a name for the dataset.
9. Under Type-Storage Profile, select one of the following options:
   - General Volume
   - Archive Volume
   - Temp Volume
10. Select a Size, either using the slider or by entering a number.
11. Select a **Block Size**.

    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**.

Snapshots

Snapshots are a read-only, point-in-time image of a dataset.

Because of their copy-on-write nature, a snapshot will initially consume no extra storage space.

A snapshot's size will grow as the files it contains change in the parent dataset or as files in the parent dataset are deleted.

Because of this, a snapshot serves as a low impact backup of a dataset and the files within it may be used to restore altered or deleted ones on the dataset.

Deleting a snapshot will release the data records it holds that are not held by the dataset or by another snapshot and return that space to the pool of available space.

Snapshot Indexing

View the existing snapshots for a dataset by selecting a dataset and clicking on the **Snapshots** tab in the window's bottom left corner.

On the left, view all snapshots that have been created in the dataset, including rolling snapshots, interval snapshots, and held snapshots.

The creation time and expiration date for each snapshot are shown with them.

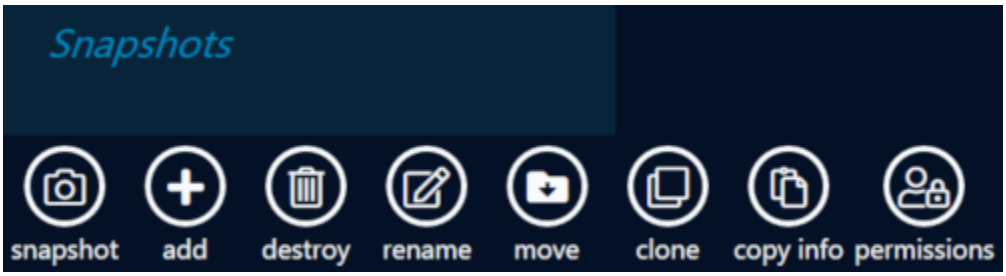

Selecting a snapshot will display information about each file present on the window's right side.

At the top, filter the snapshots to be viewed based on a range of time.

Select the **Show destroyed** box at the top to show destroyed snapshots.

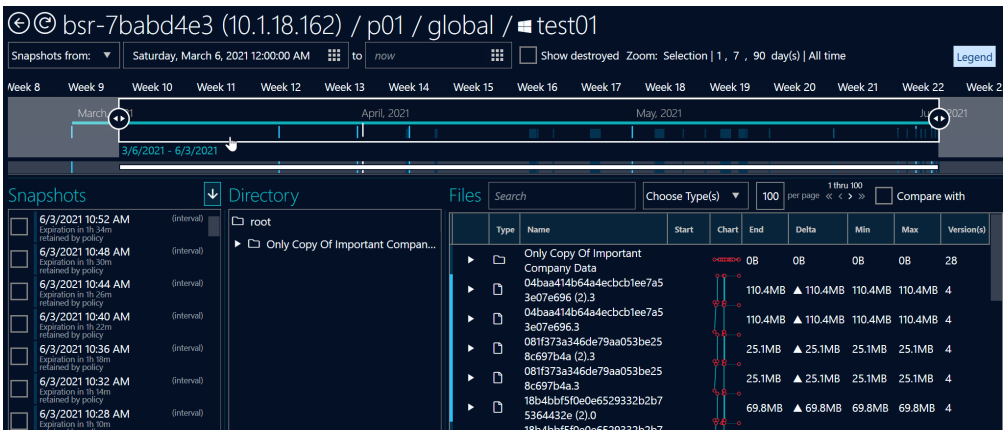

A bar spanning over a range of weeks is located at the top.

It is color-coded for indexed snapshots (light blue) and replicated snapshots (dark blue).

Hints for symbols and colors can be viewed by hovering over **Legend** in the upper right corner of the snapshots tab.

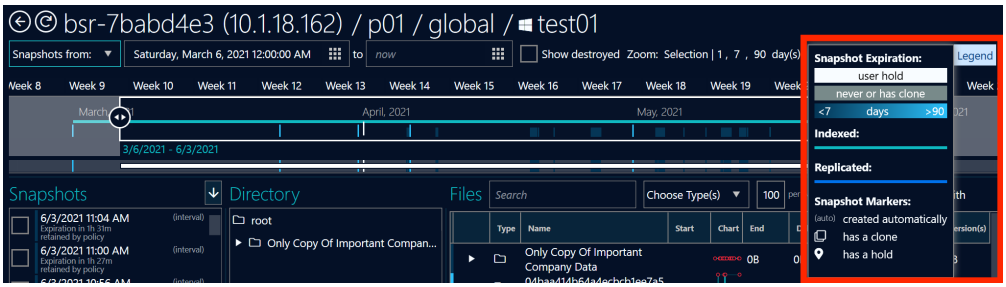

Restoring a file from a Snapshot

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**.

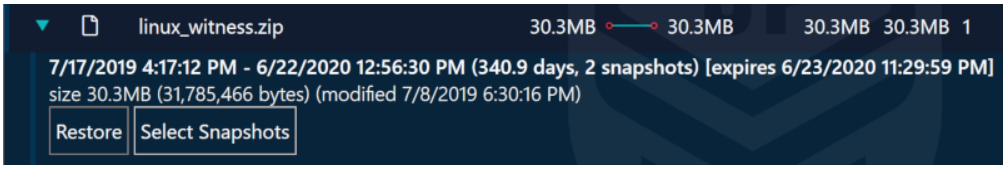

In the dialog box that shows up, choose whether the restored file should overwrite any existing file, rename the existing, or rename the restored file.

Select **Restore** to complete the action.

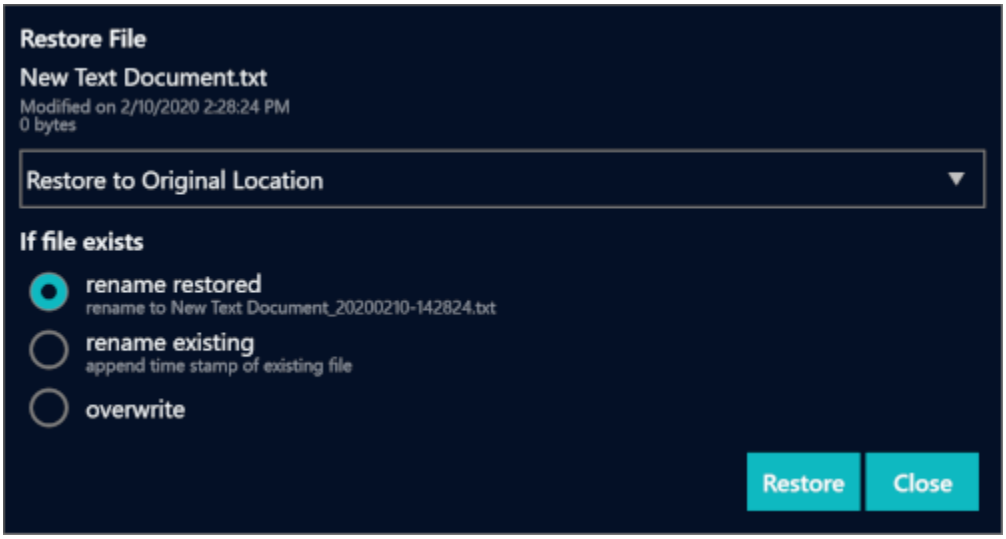

Snapshot Holds

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.

An expiration date can be set on the hold tag itself.

No snapshot will be removed from the dataset if there is a hold tag applied.

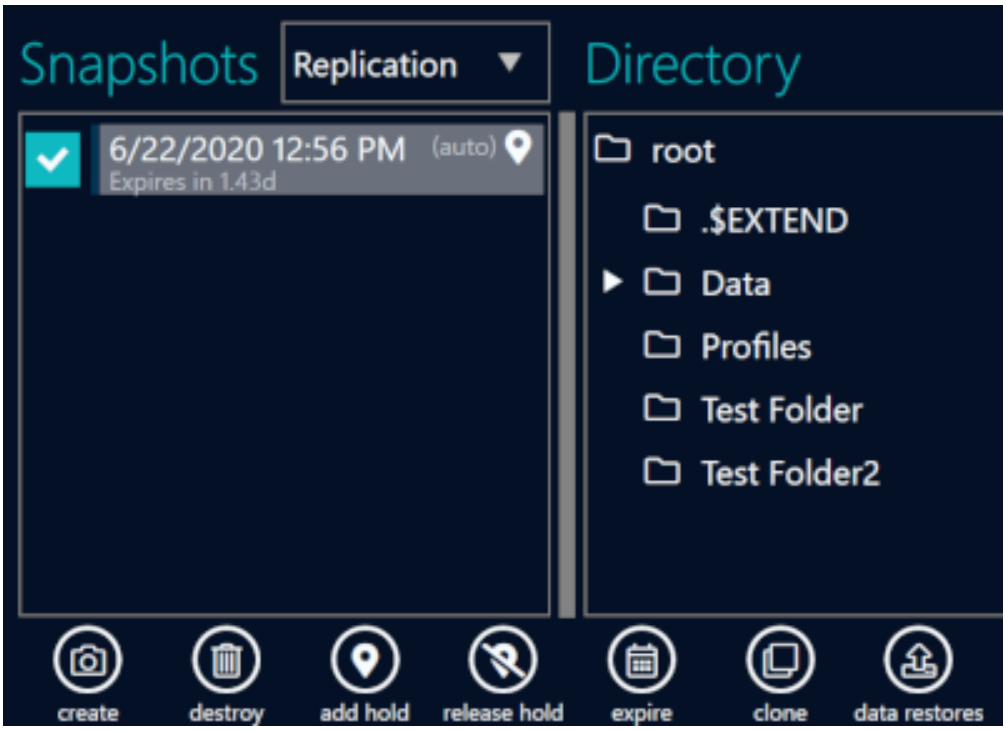

To release a hold tag you can just click release hold on the appropriate snapshots.

If a dataset is deleted, as are the snapshots that existed of the dataset.

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

Rolling Snapshots

Rolling snapshots are taken every minute and automatically expire.

There will always be five rolling snapshots.

When the sixth snapshot is created, the oldest snapshot will be deleted, always leaving five.

When an incident occurs, all snapshots are held and set to expire seven days later. The option to release the hold may be chosen after remediating the incident.

NOTE See Incidentd for more information.

Disable rolling snapshots by checking the **Prevent rolling snapshots** box.

NOTE Disabling Rolling Snapshots requires the selection to **Use Custom Protection Policy.**

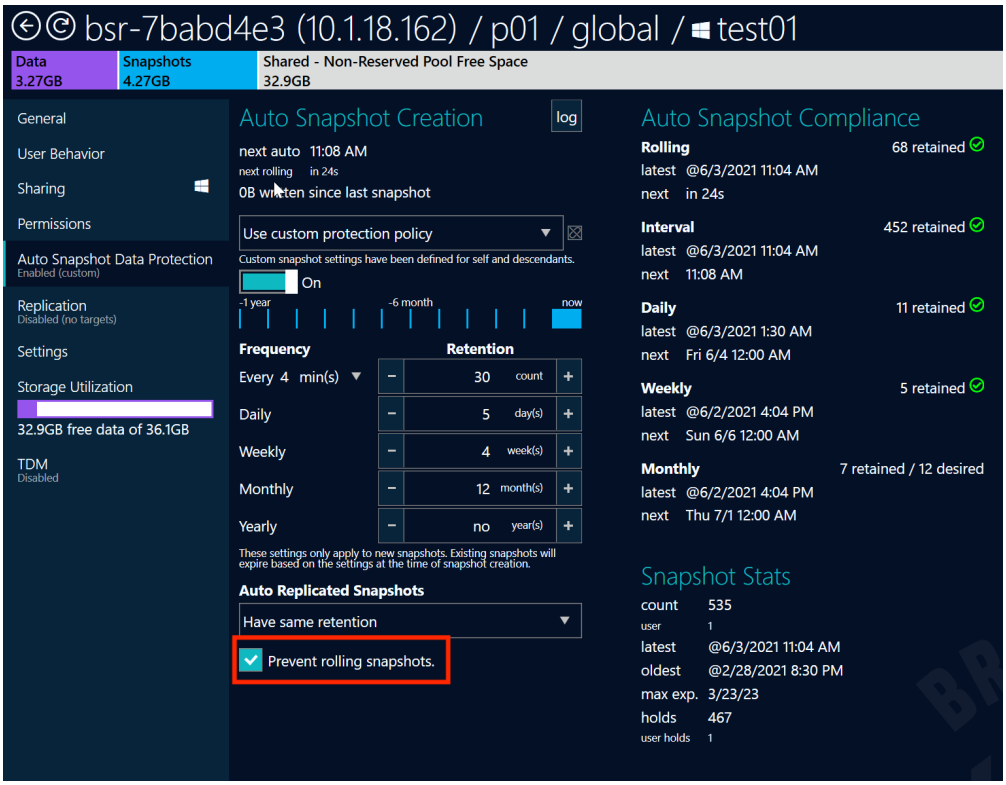

Clones

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 datasets where data is being written, snapshots should be set to run more frequently than ran 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 dataset is being replicated, keys should be exported from the local BrickStor SP and imported on the remote BrickStor SP so that the data can be recovered there.
3. Use the advanced configuration parameters to optimize replication:
    - Priorities can be set to determine which datasets will replicate first.
    - Bandwidth throttling can be configured to optimize how much bandwidth is used and at what times of day.
    - Optimize snapshot retention periods on both ends.
    - On the local system, ensure that snapshots are not aging out before they are replicated.
    - On the remote system, longer retention will consume more storage.
4. Replication peers should be on an appropriate data network that will be available and not interfere with other network traffic.
5. Setup email notifications.

Understanding Peers

BrickStor SP supports block replication between two or more pools within the same system or across systems.

To set up replication between two systems, Establish a peer relationship with the target system from the origin system.

Once the peer relationship is created 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 **Replication tab**.
3. Click on the **Add Peer** Button at the bottom left of the **Details pane**.

TIP	If peers already exist on the system, navigate to the **Add Peer** button by clicking the **+** button beside the word **Locations**.

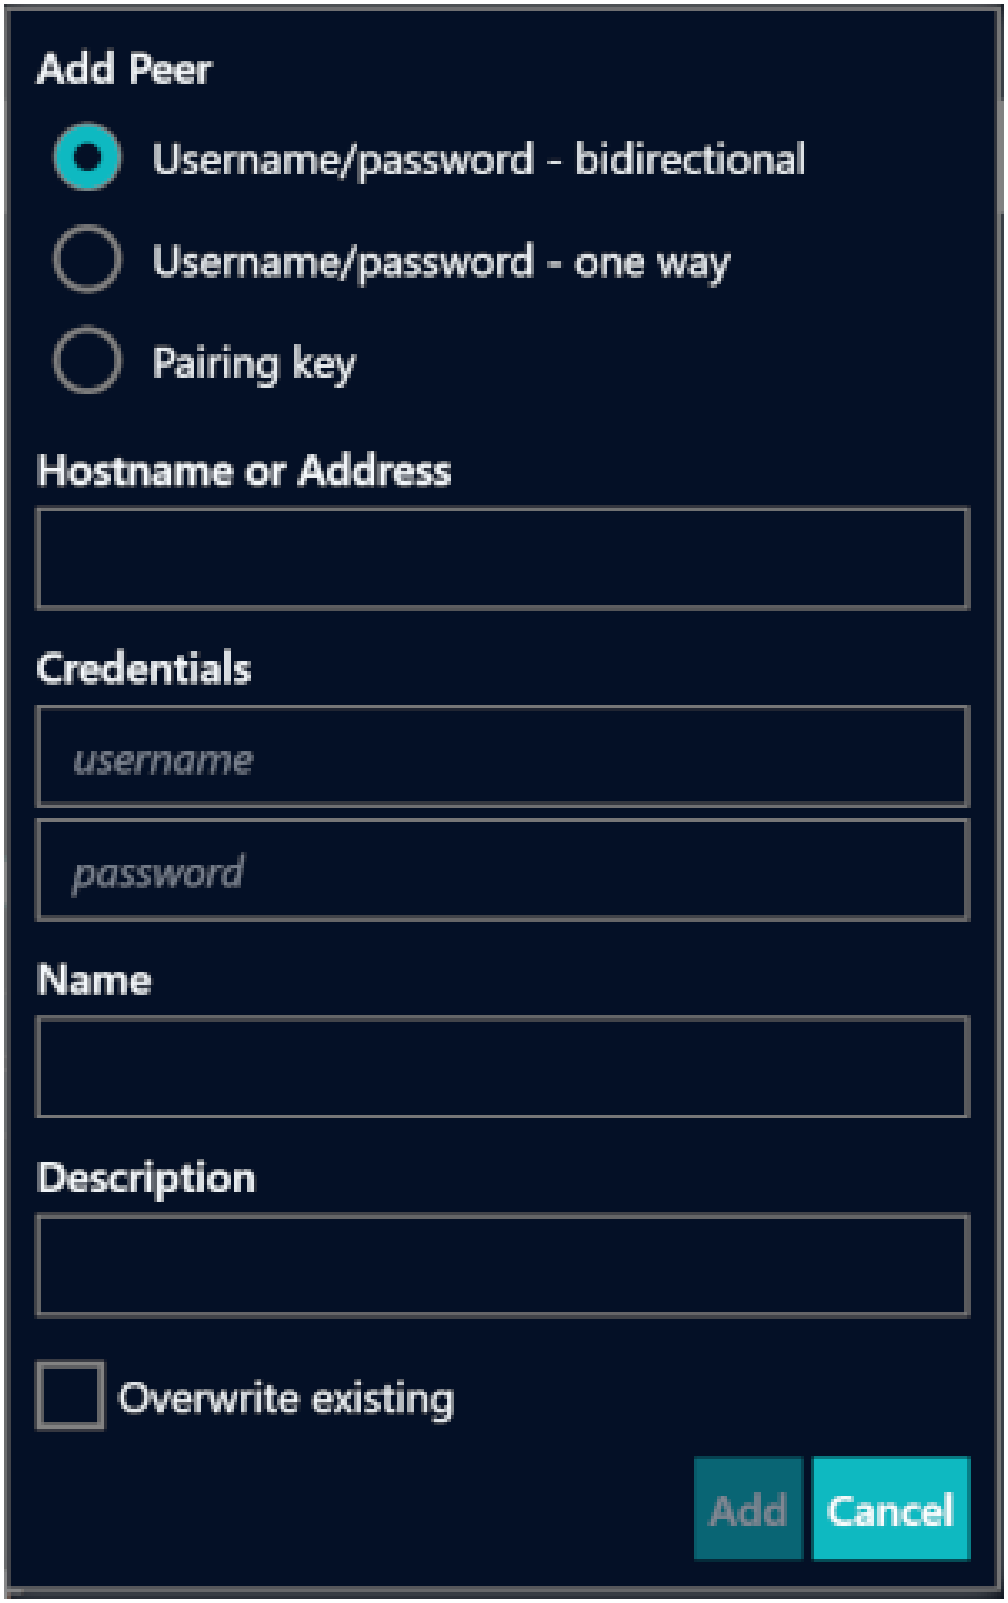

4. In the Add Peer dialog box, enter an **IP address** or **hostname** for the desired peer.

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. The BrickStor SP OS will coordinate sharing keys between the cluster nodes. If replicating to an HA cluster, ensure the use of the destination resource group's address (VNIC) in this step.

5. Enter the username and password for the desired peer.
6. Click **Add Peer**.
7. The added peer appears under the **Replication Peers** label.
8. The new peer will remain greyed out until a target has been added to that peer.
9. 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 may be encountered.

Table 2. Peer Status

This…	Means the peer is…
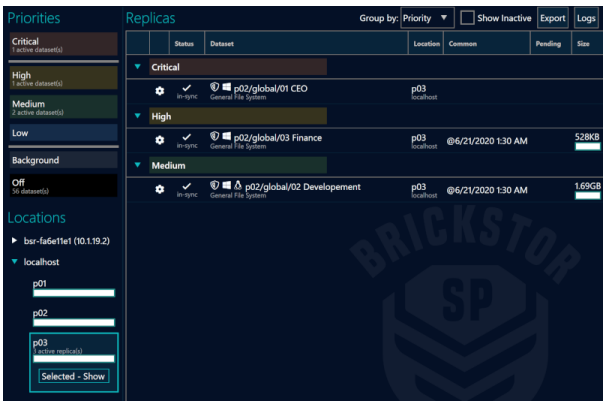	Healthy No Backlog
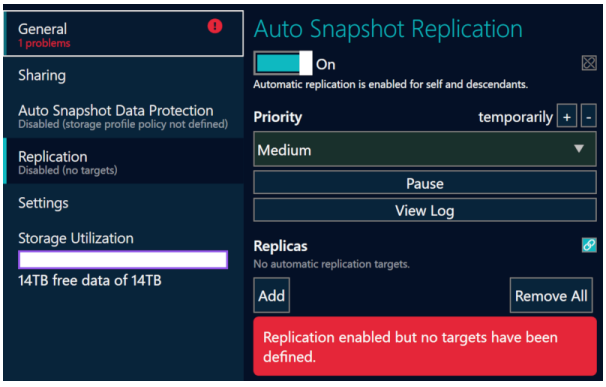	Configured without replication targets enabled for Peer

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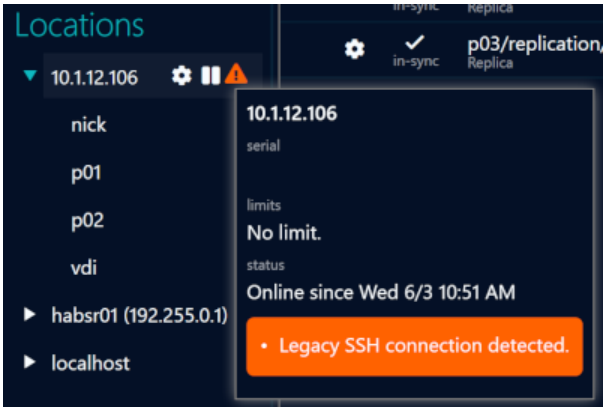	Unreachable and has a Problem, such as the target pool is not imported and will show up as [unk] or the target pool is out of space.

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.

This information is stored in file system metadata on the replication target. 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`
      - Serial number of source BrickStor
        - GUID of source dataset

Data Protection Policy Configurations

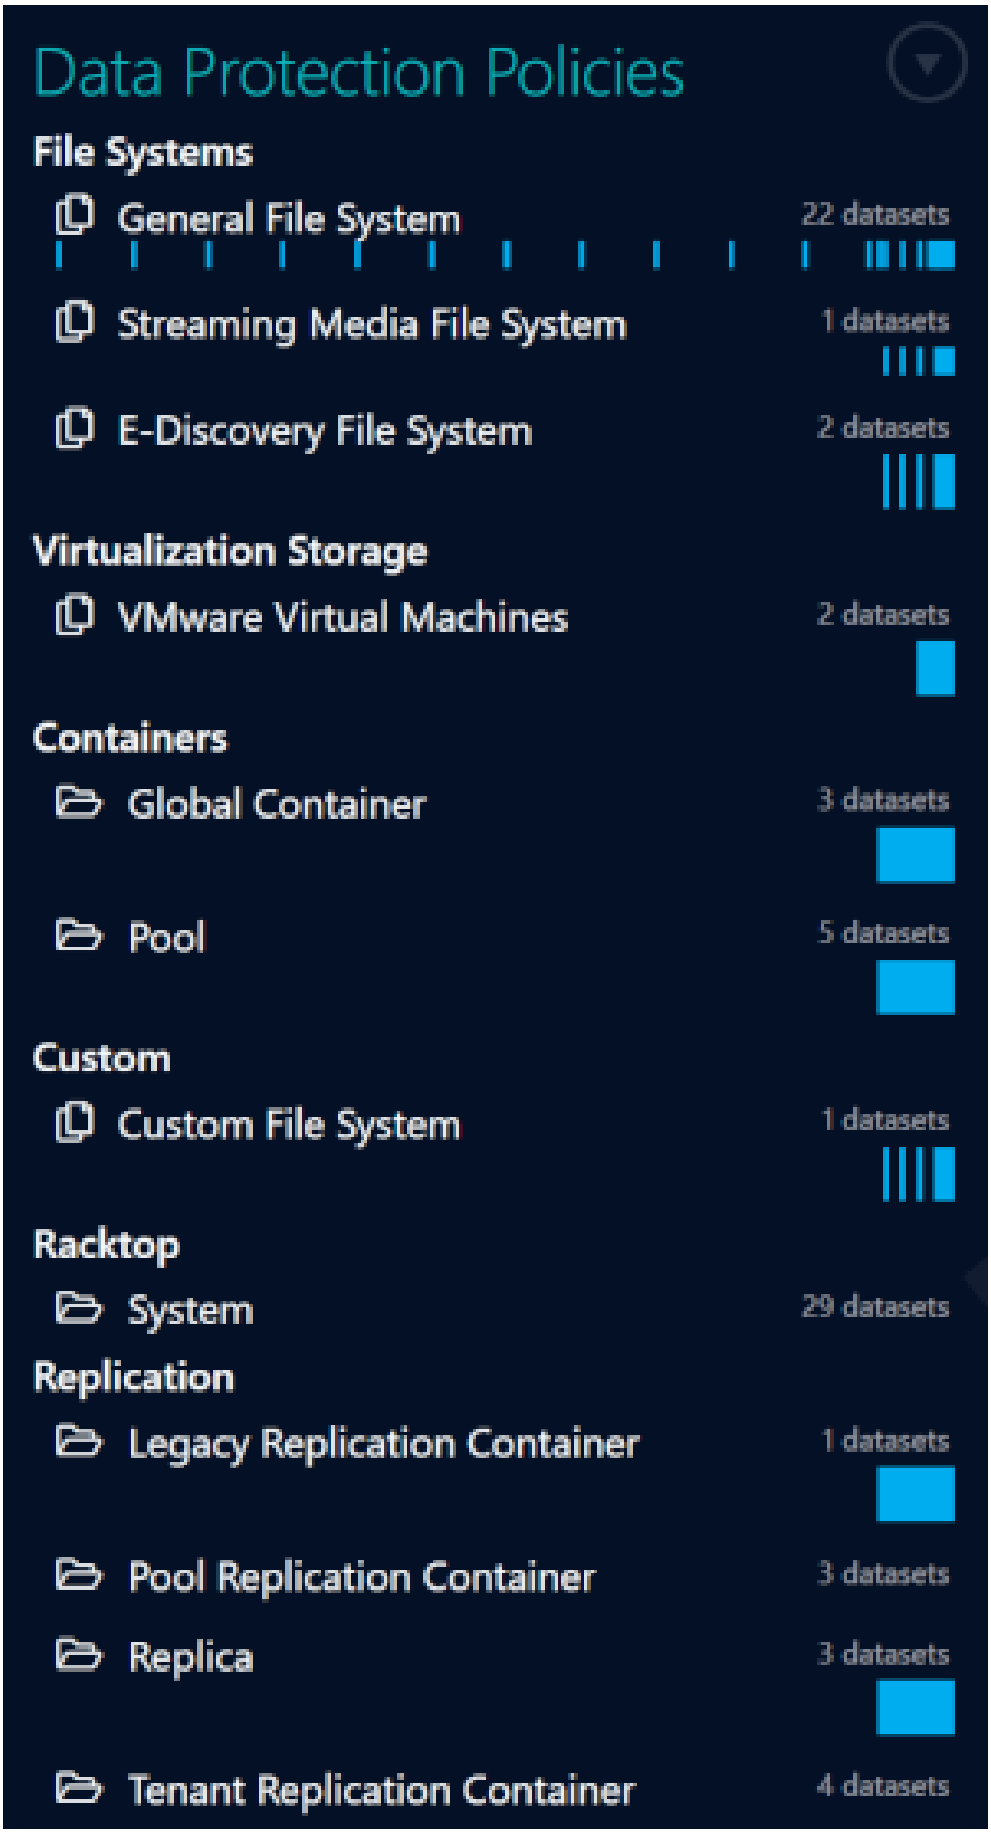
Data Protection Policies
File Systems
General File System
22 datasets
Streaming Media File System
1 datasets
E-Discovery File System
2 datasets
Virtualization Storage
VMware Virtual Machines
2 datasets
Containers
Global Container
3 datasets
Pool
5 datasets
Custom
Custom File System
1 datasets
Racktop
System
29 datasets
Replication
Legacy Replication Container
1 datasets
Pool Replication Container
3 datasets
Replica
3 datasets
Tenant Replication Container
4 datasets

Data Replication Priorities

Each replicated dataset has a priority assigned to it.

The priority determines the order that replicated datasets are sent.

The possible priorities are:

1. Critical
2. High
3. Medium
4. Low
5. Background

Critical priority datasets are always sent before datasets of any other priority.

Datasets with a priority of **Background** are always sent after any datasets of any other priority have been sent.

For **High**, **Medium**, and **Low** priority datasets, the order chosen depends on a combination of factors such as:

- The amount of data to transfer.
- The success of past replication attempts of this dataset.

The replication priority is combined with these factors to determine a 'fair' replication order to allow all datasets to make progress replicating (when possible).

Consequently, a **High** replication cannot indefinitely preempt replication of a **Medium** or **Low** priority dataset.

Likewise, a **Medium** priority dataset cannot indefintely preempt replication of a **Low** priority dataset.

Configure the Data Protection Policy for a Storage Profile

Managing Replication Details

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 to the replication details page.

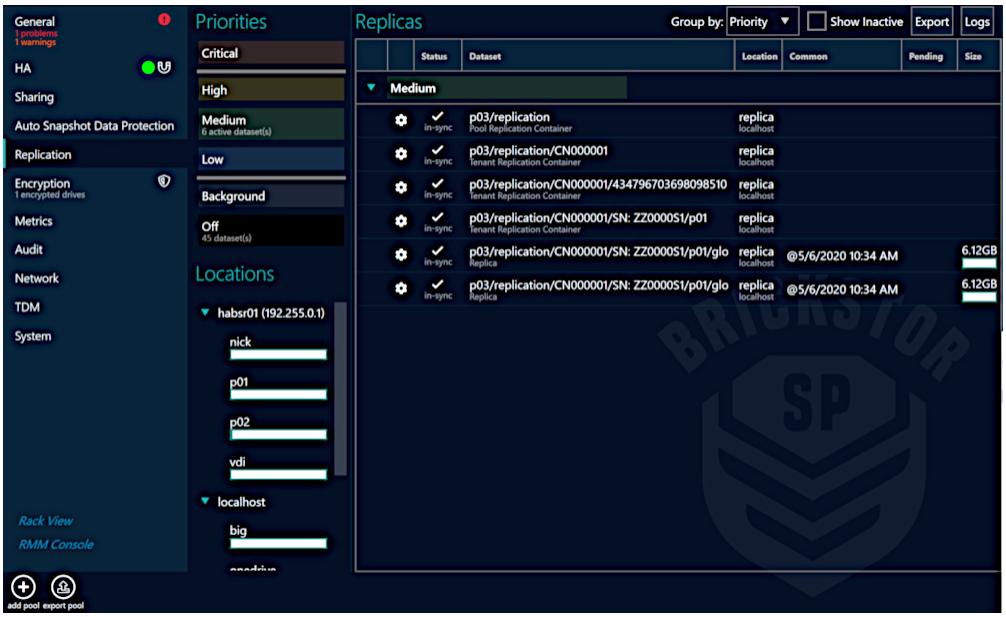

Replication Transfer History

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

Details include:

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

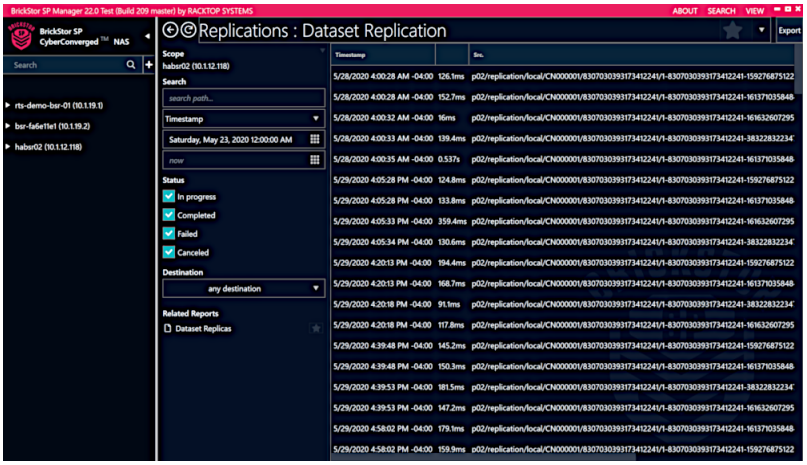

Auto Snapshot Data Protection

Within the selected dataset, click on the **Auto Snapshot Data Protection** tab.

Set a custom profile protection policy under the Auto Snapshot Creation section and filter as needed.

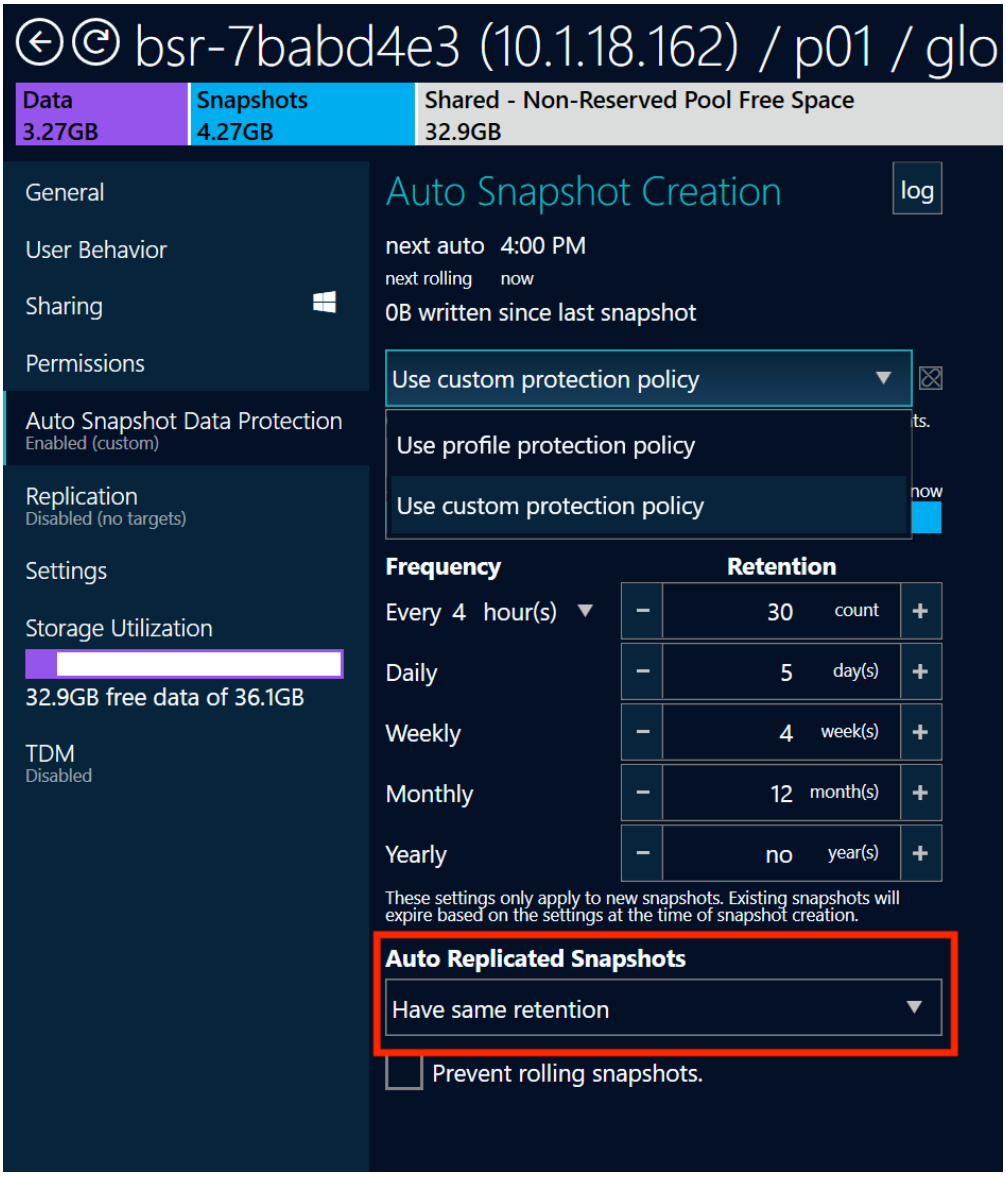

Choose whether to have the same or alternate retention under Auto Replicated Snapshots.

To the right (the Auto Snapshot Compliance area(), includes the number of snapshots retained and desired, as well as the latest snapshot and next snapshot time for all rolling, interval, weekly, monthly, and yearly snapshots.

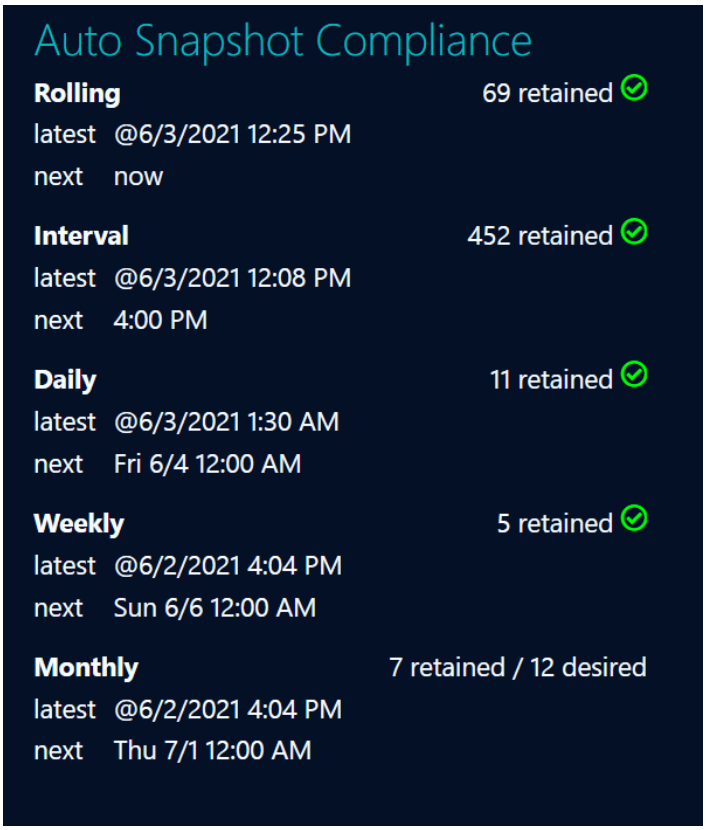

The snapshot stats display shows the count, users, latest and oldest snapshots, max expiration, holds, and user holds.

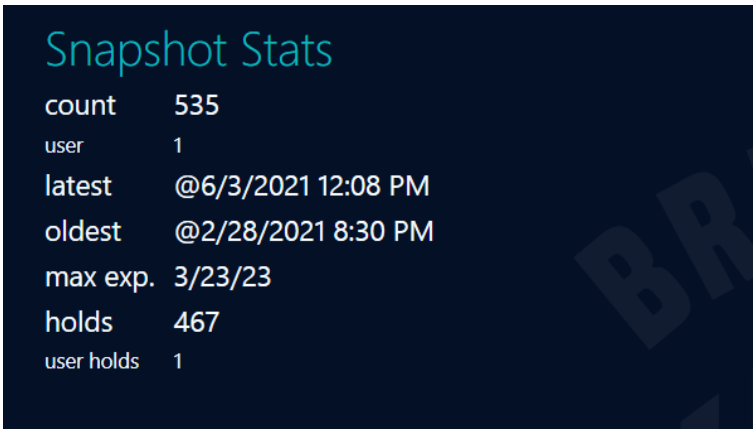

Further to the right (the Snapshot Indexing area), displays the following information and allows the user to toggle the on and off switch for indexing snapshots.

Snapshot Indexing also gives the option to regenerate the index, which will prompt the user with a time consumption warning.

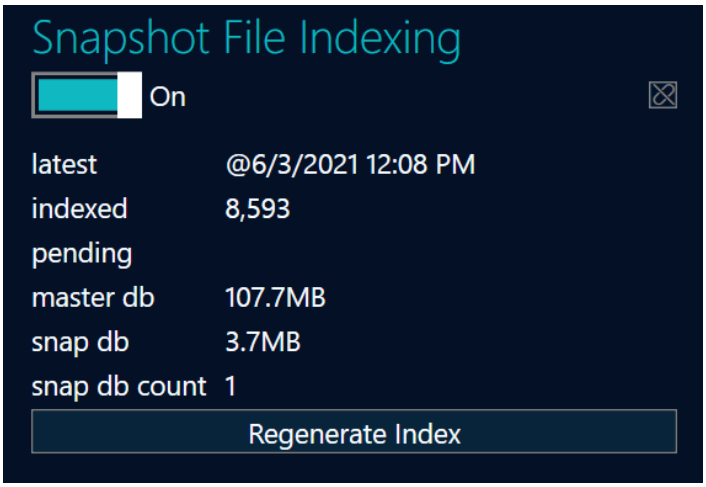

Further to the right under reports, click **Auto Snapshot Creation: Policies**.

Here, set the minimum and maximum policy by selecting them with the toggle button.

Once selected, the user can filter and add the needed specifications.

NOTE There may be a prompted alert if too many or too few snapshots are selected.

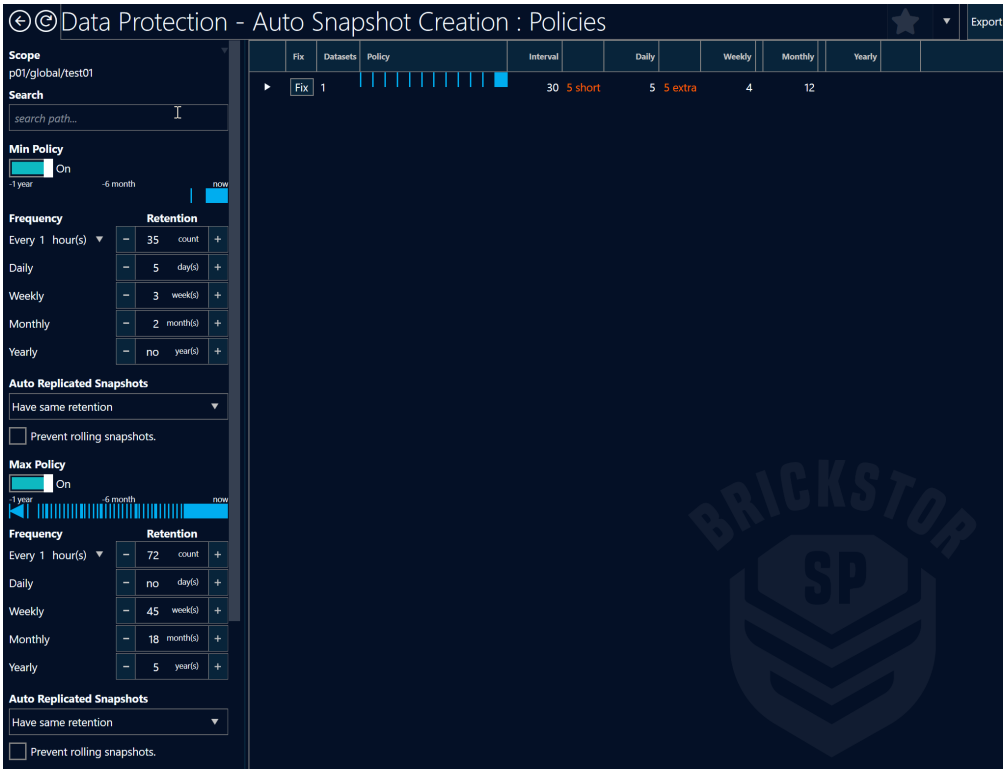

User Behavior Auditing and Analysis

User Behavior Auditing allows the ability to track how end users interact with data stored on Brickstor SP.

User Behavior logs the operations for each file made by applications and users, such as file creation, movement, deletions, etc.

BrickStor displays this information in real-time reports and graphs.

Enable User Behavior at the pool level or the dataset level:

- BrickStor SP logs the behavior of users at the system level where it was configured and its descendants.
  - For example, if User Behavior at the Pool Level is enabled, 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 **Sharing** tab.
3. Under User Behavior, click the toggle button to **On**.

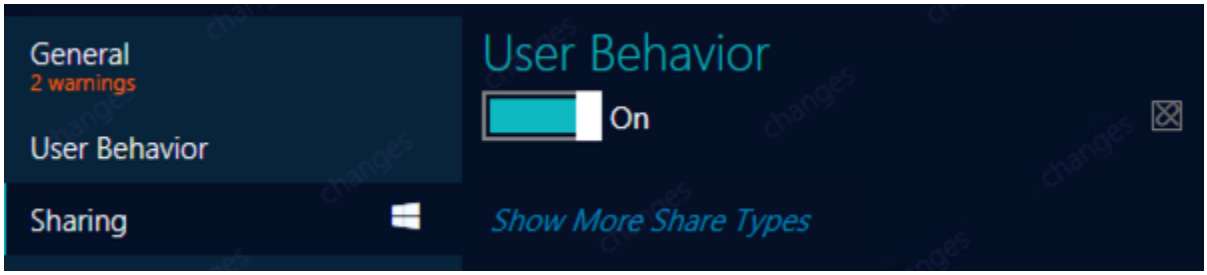

4. In the Changes pane, click **Commit Changes**.

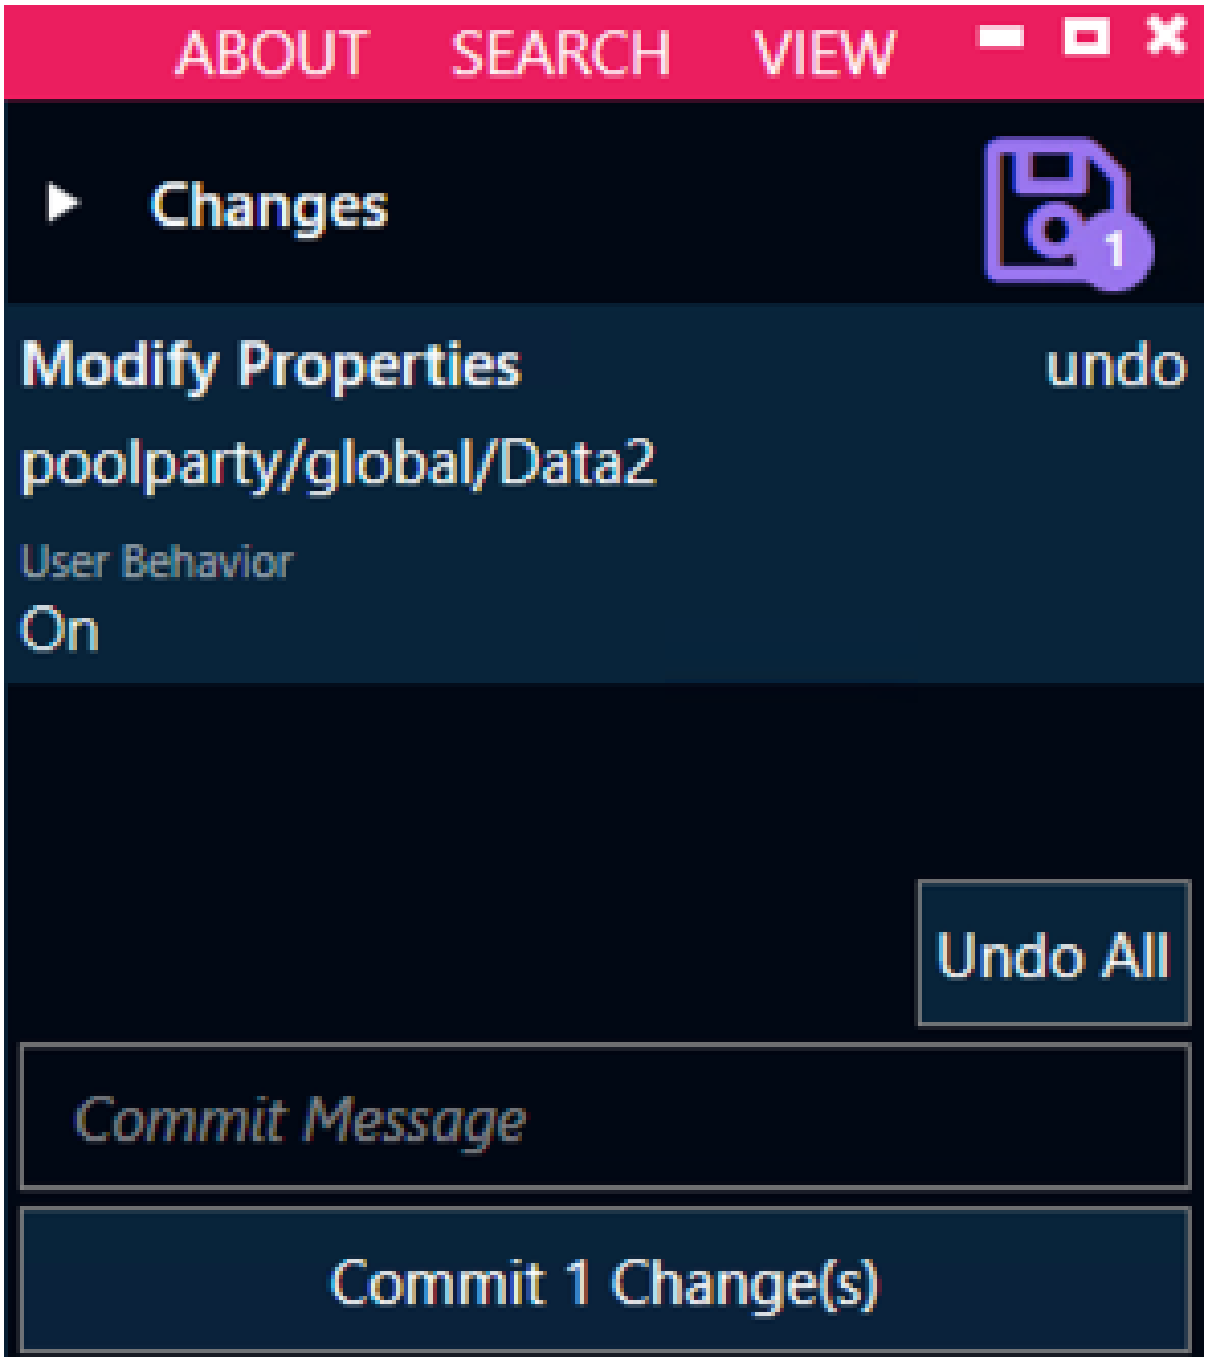

User Behavior Audit

After User Behavior is enabled, BrickStor SP displays an overview of all user actions initiated from that point.

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.

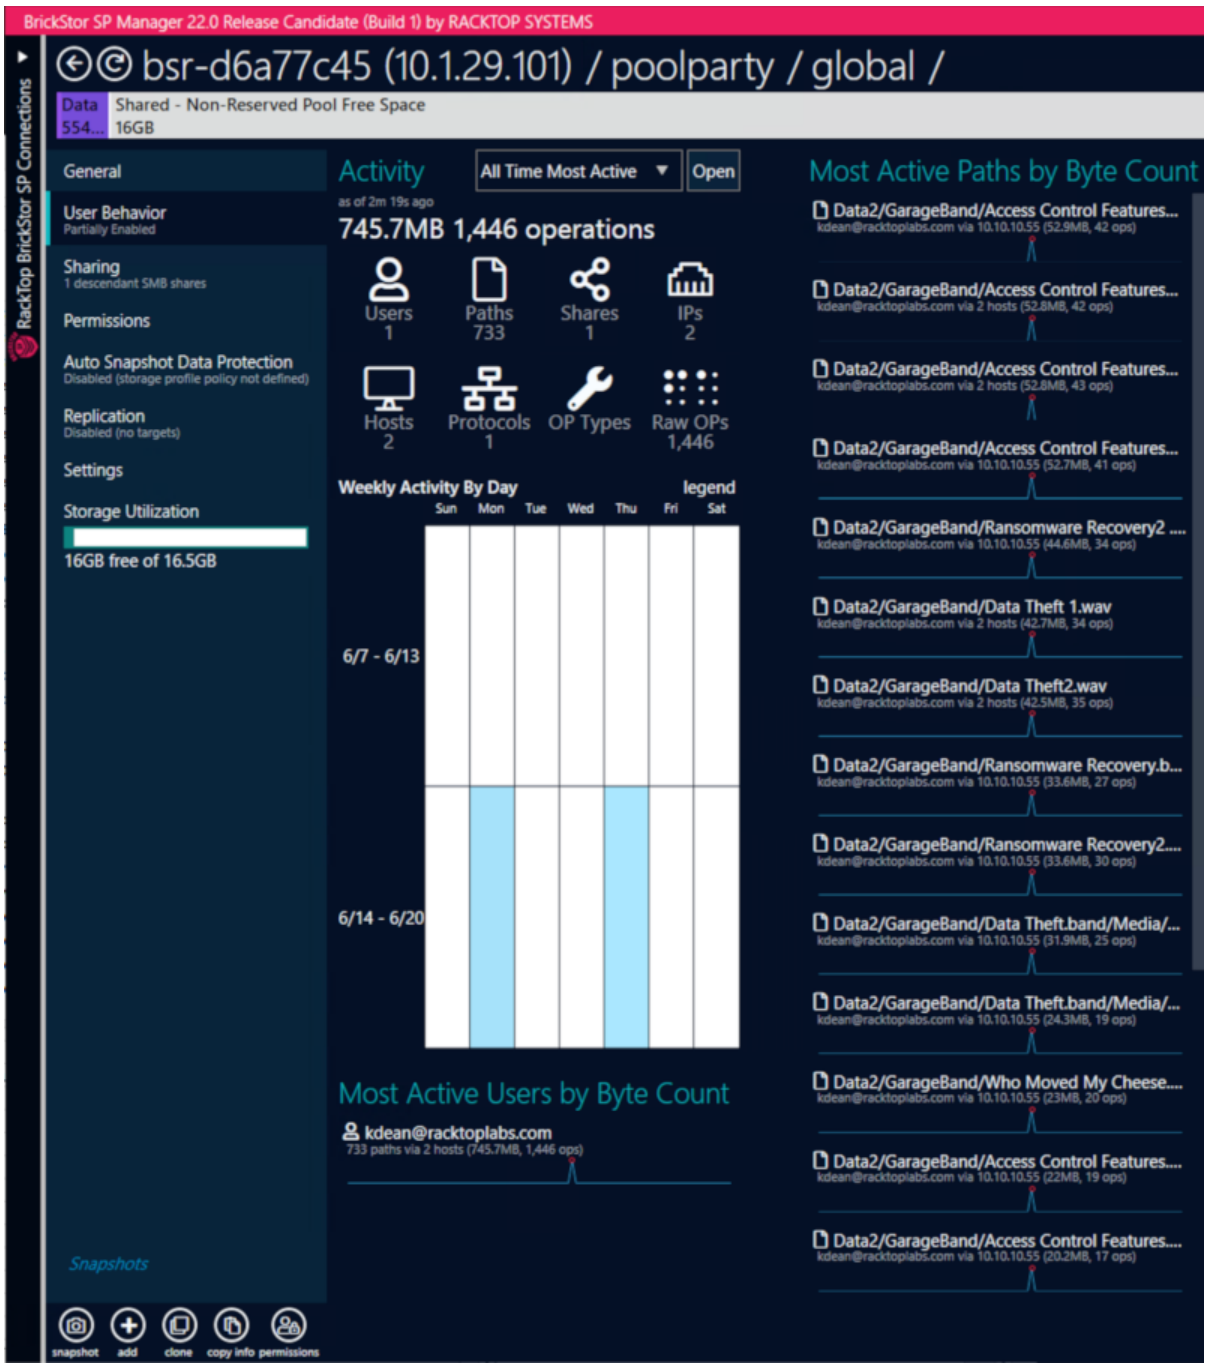

Most of the content here can be clicked on and will lead to the *Activity* page.

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, begin by running `setup`.

```
root@bsr-3841af53:~# setup
RackTop Cyberconverged NAS
Setup Utility
Copyright 2022 RackTop Systems, Inc.

 Main Menu

 1. Configure RMM interface.
  2. Configure nodename.
  3. Configure network interface.
  4. Configure aggregate network interface.
  5. Configure NTP settings.
  6. Configure DNS settings.
  7. Disable system service connections to the Internet.
  8. Configure TimeZone.
  9. Restart appliance.
10. System Information and Administration.
11. Exit Setup Utility.

Select menu option and press enter or press enter to exit.
```

- Select **Option 10**, and press **Enter**.

NOTE Use CTRL-C to exit at anytime.

```
RackTop Cyberconverged NAS
Setup Utility
Copyright 2022 RackTop Systems, Inc.

 System Information and Administration Menu

 1. Operating System Version.
  2. Hardware list.
  3. Additional System Information
  4. License Information
  5. Show interface links.
  6. Change local password.
  7. Add local User account.
  8. Remove local User account.
  9. Review current state of services.
  10. Enable or disable service.
  11. Add system to Active Directory.
  12. Check Active Directory.
  13. IO Status Check.
  14. Configure Syslog Forwarding.
  15. Add a license key to system.
  16. Upgrade operating system.
  17. Support Bundle.
  18. Start a shell.

 Please select menu option and press enter or press enter to return to main menu.
```

- In the System Information and Administration Menu, select **14 - Configure Syslog Forwarding**, and press **Enter**.

```
1. Syslog forwarding.
  2. UB forwarding.
  3. Disable Syslog forwarding.
  4. Disable UB forwarding.

 Select option above:
```

- Select **2 - UB forwarding**, and press **Enter**.

```
What protocol would you like to use: (options: tcp/udp)
```

- Enter tcp, and press **Enter**.

```
  What is the IP Addess to the server that you would like to send to:
```

- Enter <IP of Your BrickStor SP>, and press **Enter**.

Active Defense

Active Defense is the BrickStor SP feature that detects ransomware attacks, malware activity, and other types of unusual activity on file systems in real time.

When a Rule is triggered by suspicious activity, an Incident is created. This will trigger an alert, as well as initiate any of several actions such as blocking the user or IP address from which the attack originates. The creation of an Incident also causes Data Protection to create a point-in-time read-only snapshot of the affected file system to aid in isolation and recovery of affected files. Once an Incident is generated, an administrator may acknowledge it and remove any blocks that were put in place.

Active Defense is managed using the Security Incidents screen of BrickStor SP Manager.

Security Incident Display and Workflows

To show Security Incidents:

1. In BrickStor SP Manager, navigate to the **General** tab on the managed appliance.
2. Click **Security Incidents** near the bottom of the Details pane.

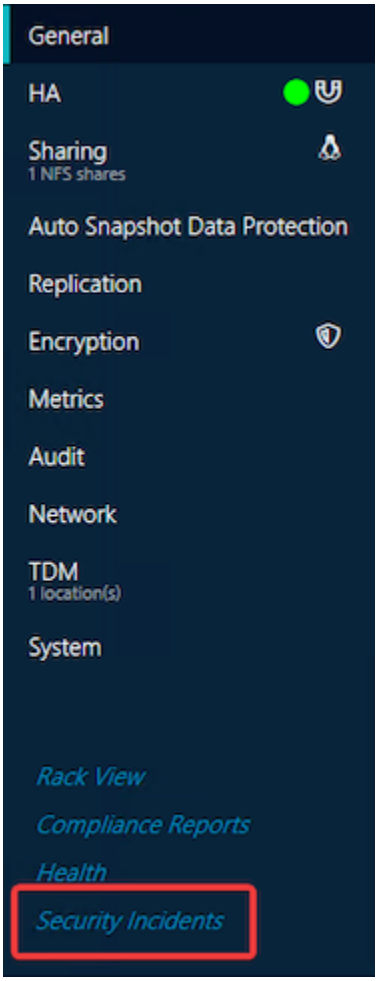

Incidents will be listed in the Security Incidents section with information including the type of incident,

user, endpoint IP address, and timestamp. Use filters to sort the incidents by date/time. Selecting the **Closed** checkbox will show incidents that have been closed.

Selecting an incident will show additional information and buttons to acknowledge or remediate the incident and provide actions to add watchers, notes, and more.

Incident Details

- **Type** - type of incident.
- **Score** - severity score 0-10 (0-1.0, with 1.0 being absolutely confident).
- **User** - user login which triggered the incident.
- **IP** - user endpoint incident origination.
- **Created** - incident creation date and time.
- **Acknowledged** - which administrator acknowledged the incident and when (date/time).
- **Closed** - which administrator remediated the incident and when (date/time).

Actions

The **Actions** section displays the actions that were taken in reaction to the incident. The status balloon next to each action indicates the action's status. **Green** corresponds to the action currently being enforced. **Grey** indicates the action has been lifted by the system administrator.

The **Lift** button allow a system administrator to remediate the incident by Lifting or unblocking the restrictions created by the incident.

- **Block Host** - client endpoint IP address is blocked from accessing the shares.
- **Block User** - authenticated user login is blocked from accessing the shares.
- **Hold Snapshots** - related snapshots are held and their expiration time is extended.
- **Prevent Auto Reapply** - This drop-down allows the administrator to choose to create a time-limited exception for the user account, IP, and specific incident type.

Datasets

This section will show all datasets affected by the selected incident along with each dataset's **Activity** and **Snaps** buttons. Clicking the **Activity** button will open the User Behavior management screen filtering view to show activity related to this dataset. **Snaps** will open the Snapshot management screen of Data Protection.

Watchers

Watchers can be added to the incident in order to receive emails about the attack. This is done by selecting the **Add** icon next to Watchers and adding the email address of the user. Lift or reapply the actions of blocking the user, IP, and holding snaps by selecting Lift and checking off which action to lift/reapply.

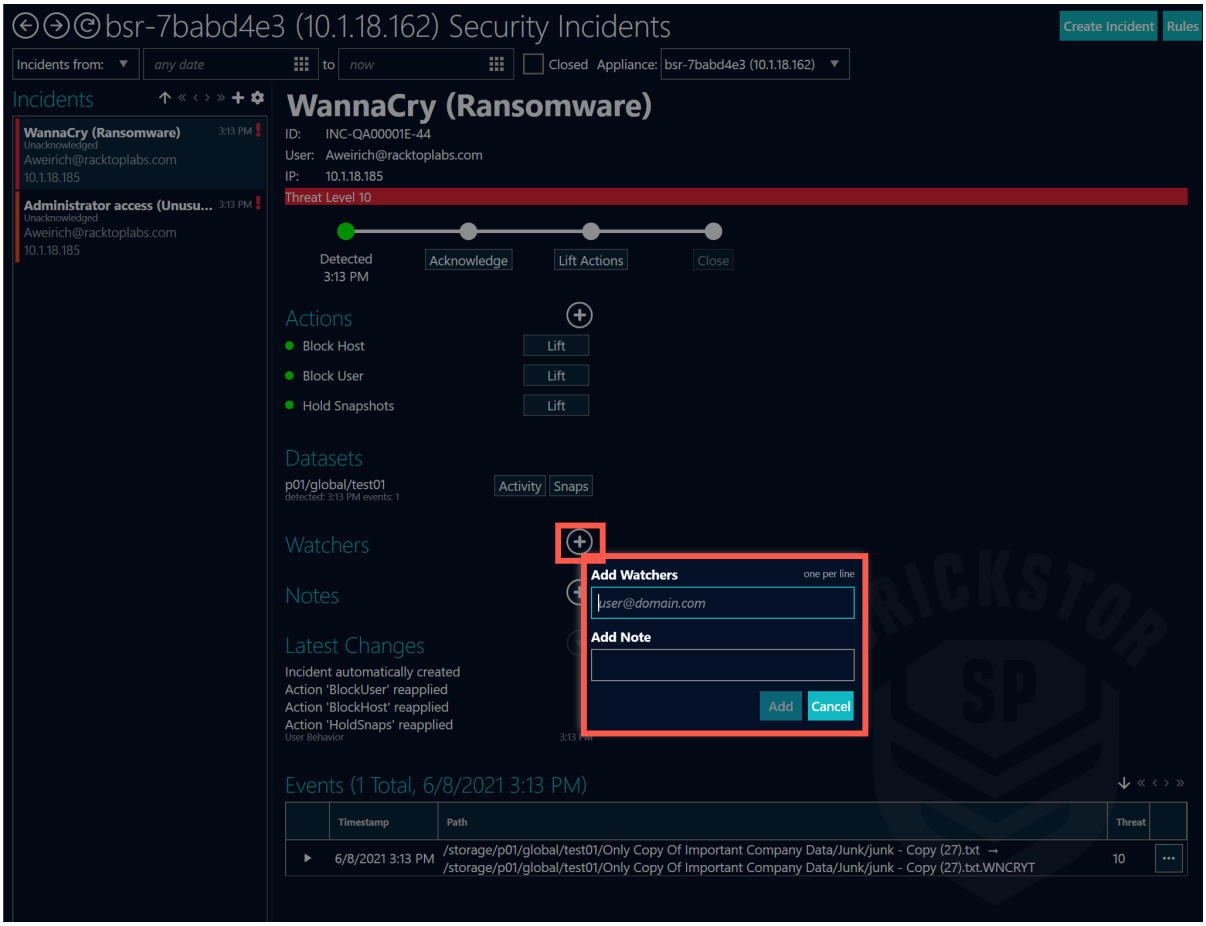

Notes

The Notes section will list any notes added to the incident. Notes can be added, edited or deleted at any point until an incident is Closed. It is also possible to add a note while adding watchers.

To append a note:

1. Click the plus (**+**) icon next to **Notes**
2. Enter message text
3. Click the **Add Note** button to save it

Recent Changes

The Recent Changes section shows audit log events associated with this incident starting from when it was first detected.

Events

The Events section lists all events triggered by the user activity for this incident.

Manual Incident Creation

It is possible to manually create incidents and to apply actions to or alert on the incident.

Press the **Create Incident** button to open a the incident details menu.

Update the fields for the incident category, name, assigned threat level, involved user, dataset, host, and any notes regarding the incident.

Watchers and actions may be assigned to the incident to block the user or host from access and alert on any occurrences of such access being attempted.

Manual Rule Creation

You have the option to create a rule in the incidents tab. Manual rule creation allows you to add the category of the incident, score, user, host, datasets, watchers, and apply any actions as well as create a custom action. It also allows you to define the rule type (continue processing rules, stop processing rules, or do not open incident).

There is also the option of adding an expiration date on the rule. You can do this by clicking rules in the upper right hand corner of the incidents tab, clicking edit at the bottom of the **Rules** tab, and then clicking **Add**.

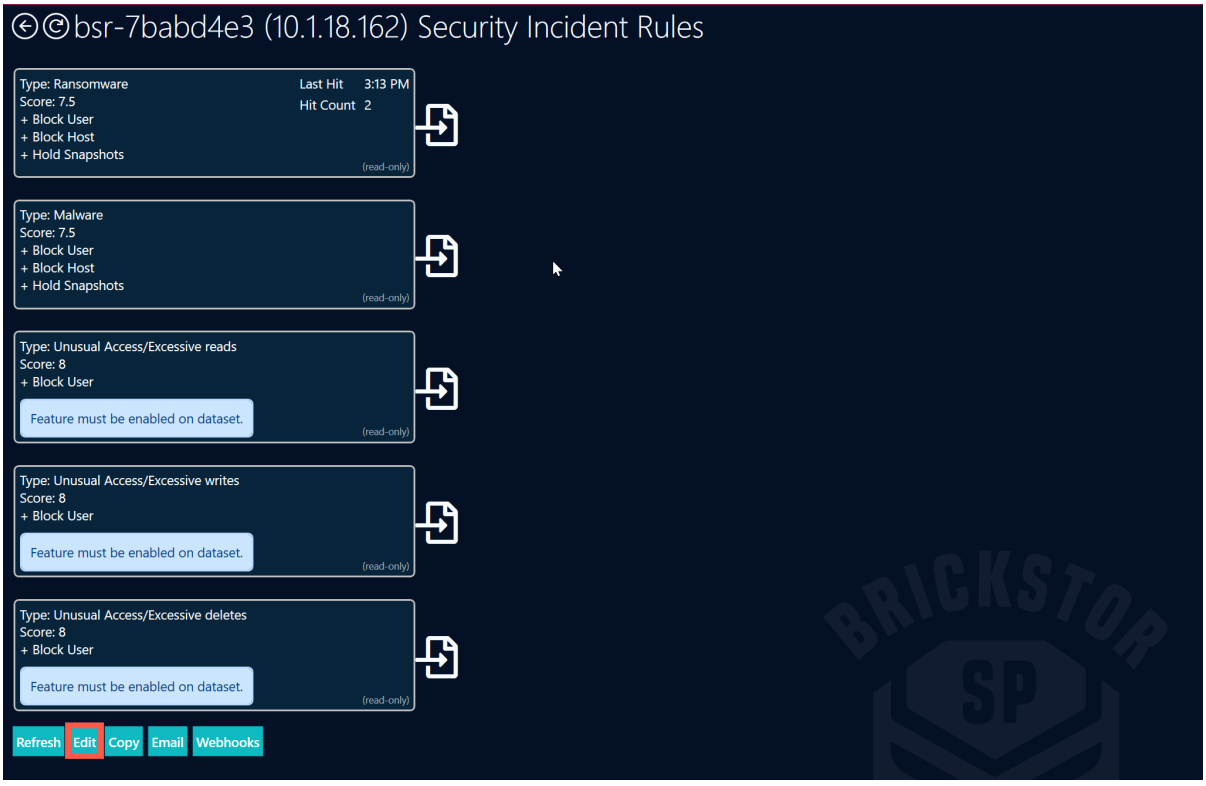

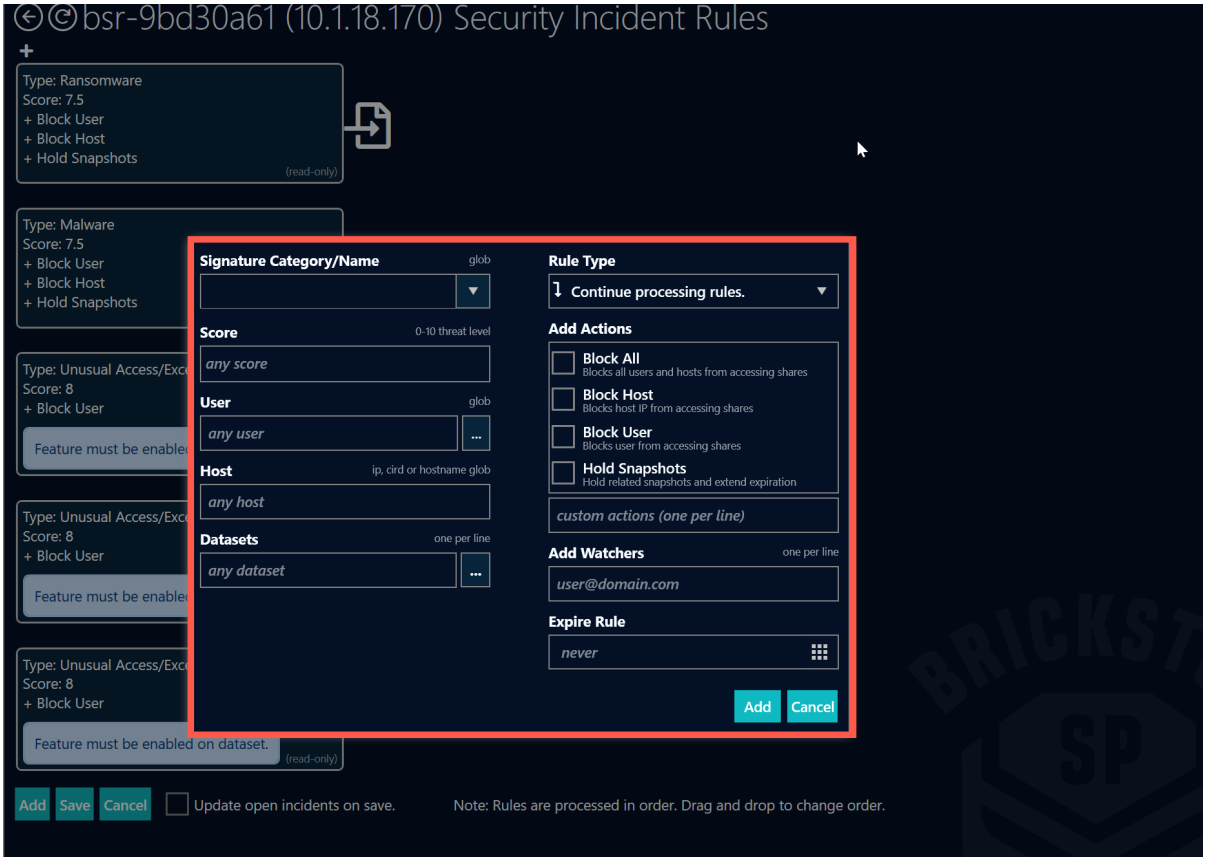

NOTE

The user has the option to access the ***security incident rules*** through the ***system*** tab at the appliance level.

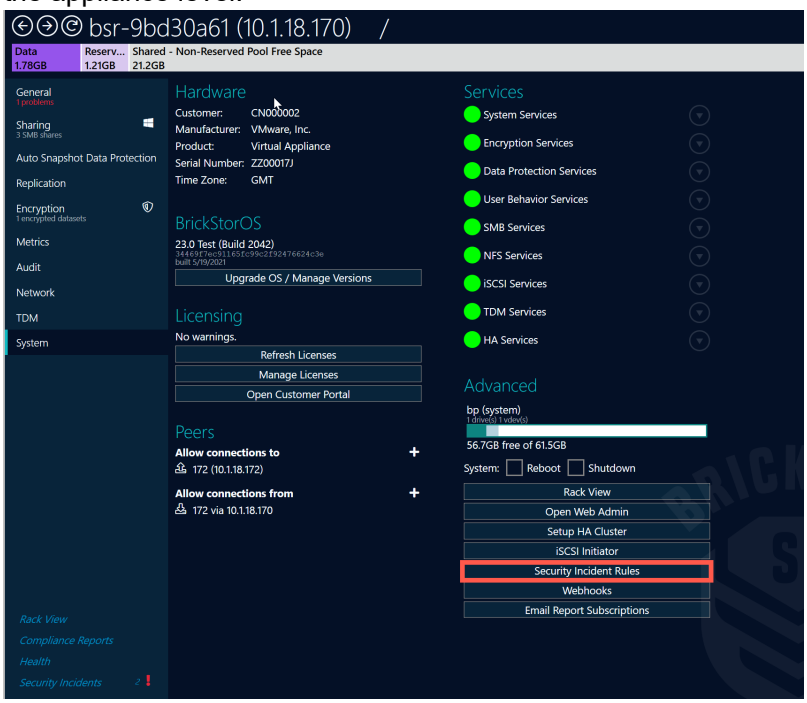

Assessors and Rules

Assessors and rules are used by Active Defense to constantly analyze the activity of the system or datasets. Any activity that matches the criteria set forth in each rule or assessor causes an Incident to be created with predetermined actions and alerts activated.

The list of Assessors and Rules can be viewed by clicking the **Rules** button on the **Security Incidents** screen.

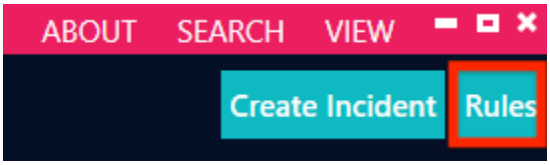

Assessors

Assessors include the following:

- Ransomware Protection
- Malware Protection
- Unusually high read activity
- Unusually high write activity
- Unusually high delete activity
- Administrator write activity
- Administrator delete activity

Assessors activity on the system. The rules of these assessors are explained further in the following sections.

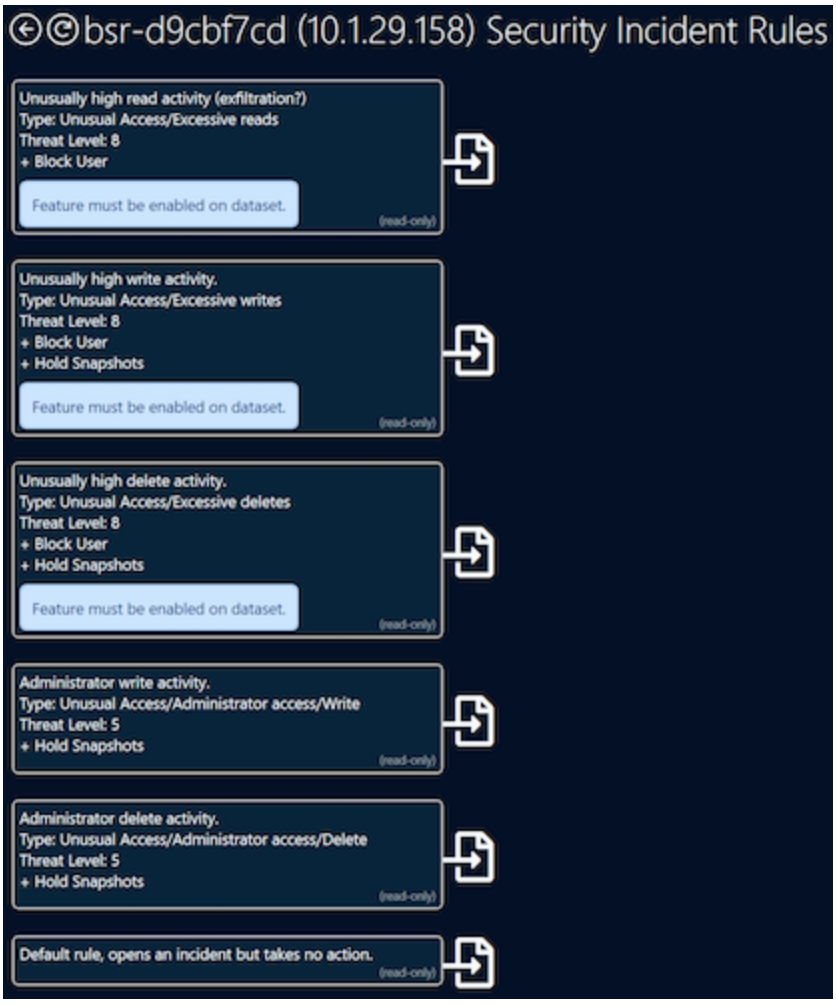
bsr-d9cbf7cd (10.1.29.158) Security Incident Rules
Unusually high read activity (exfiltration?)
Type: Unusual Access/Excessive reads
Threat Level: 8
+ Block User
Feature must be enabled on dataset.
(read-only)
Unusually high write activity.
Type: Unusual Access/Excessive writes
Threat Level: 8
+ Block User
+ Hold Snapshots
Feature must be enabled on dataset.
(read-only)
Unusually high delete activity.
Type: Unusual Access/Excessive deletes
Threat Level: 8
+ Block User
+ Hold Snapshots
Feature must be enabled on dataset.
(read-only)
Administrator write activity.
Type: Unusual Access/Administrator access/Write
Threat Level: 5
+ Hold Snapshots
(read-only)
Administrator delete activity.
Type: Unusual Access/Administrator access/Delete
Threat Level: 5
+ Hold Snapshots
(read-only)
Default rule, opens an incident but takes no action.
(read-only)

Ransomware & Malware Protection

BrickStor SP, when detecting a potential ransomware or malware attack will immediately ban the suspected agent, and place recent snapshots on hold so that they may be reinstated if needed. Moreover, BrickStor SP will provide detailed information of the agent, time of attack, and threatened files.

Insider Threat

The Excessive File Access feature is a part of BrickStor SP's Active Defense capabilities, and has the ability to detect various excessive file operations including reads, writes and delete operations.

The option to enable Excessive File Access is on the **sharing tab** for a dataset and is configurable per dataset.

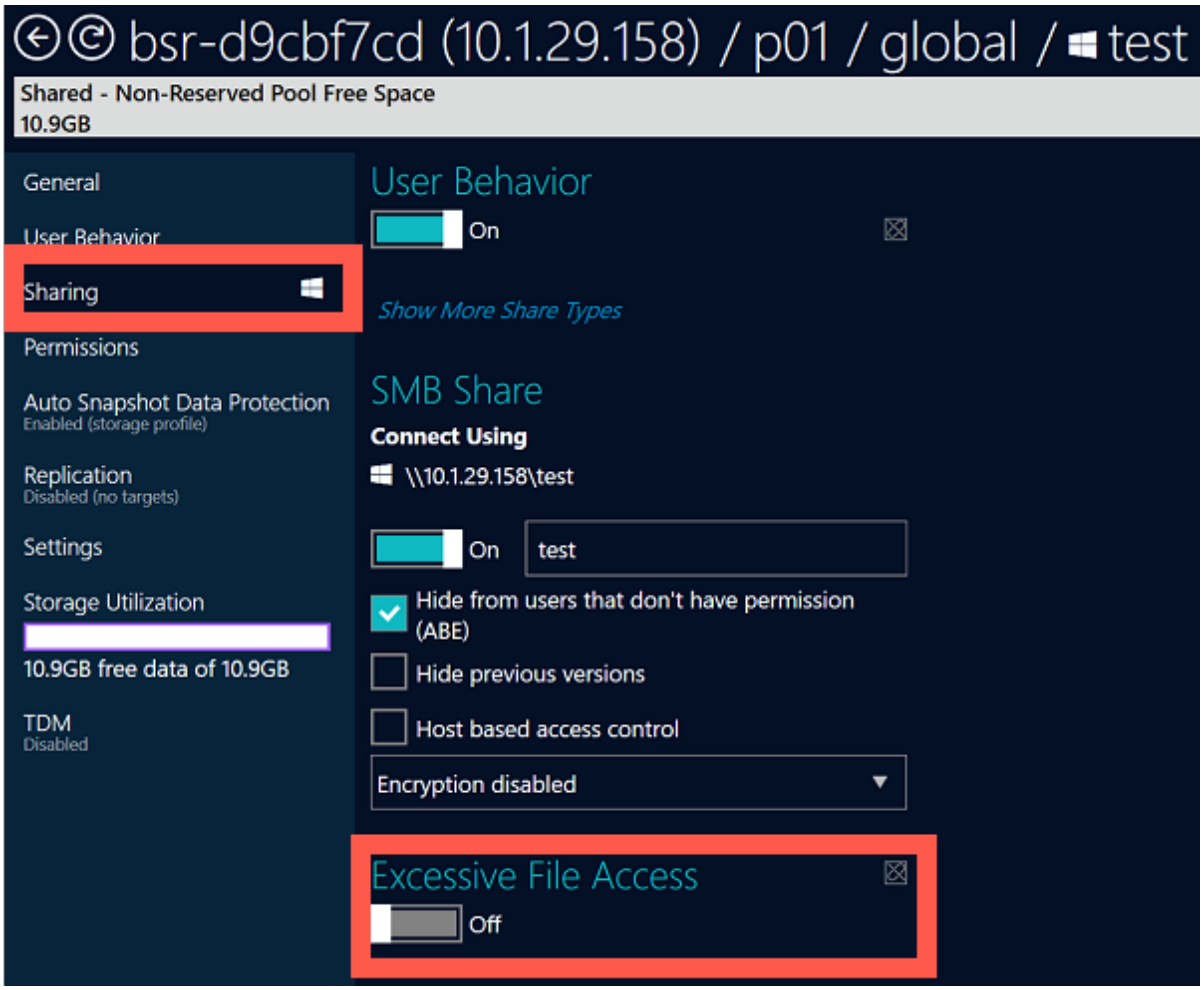

Enabling the **Excessive File Access** option will open a new dialogue box that allows the configuration for how many file operations to track per minute.

For each of the three file operation trackers, there are options to **Notify After** and **Block After**.

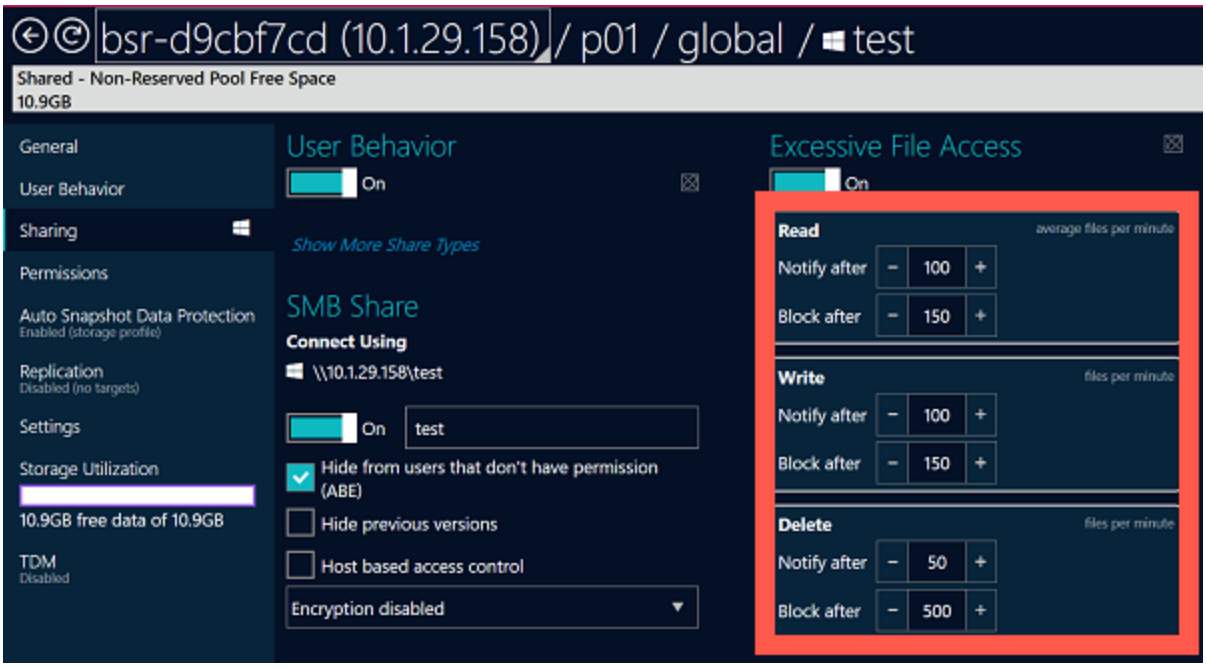

Once the **Notify After** threshold for a certain file operation has been reached, an incident will be created which can be viewed on the **Security Incidents** screen.

This will display the type of incident that has occurred, the user, the host IP for where the activity came from, and the dataset that was affected.

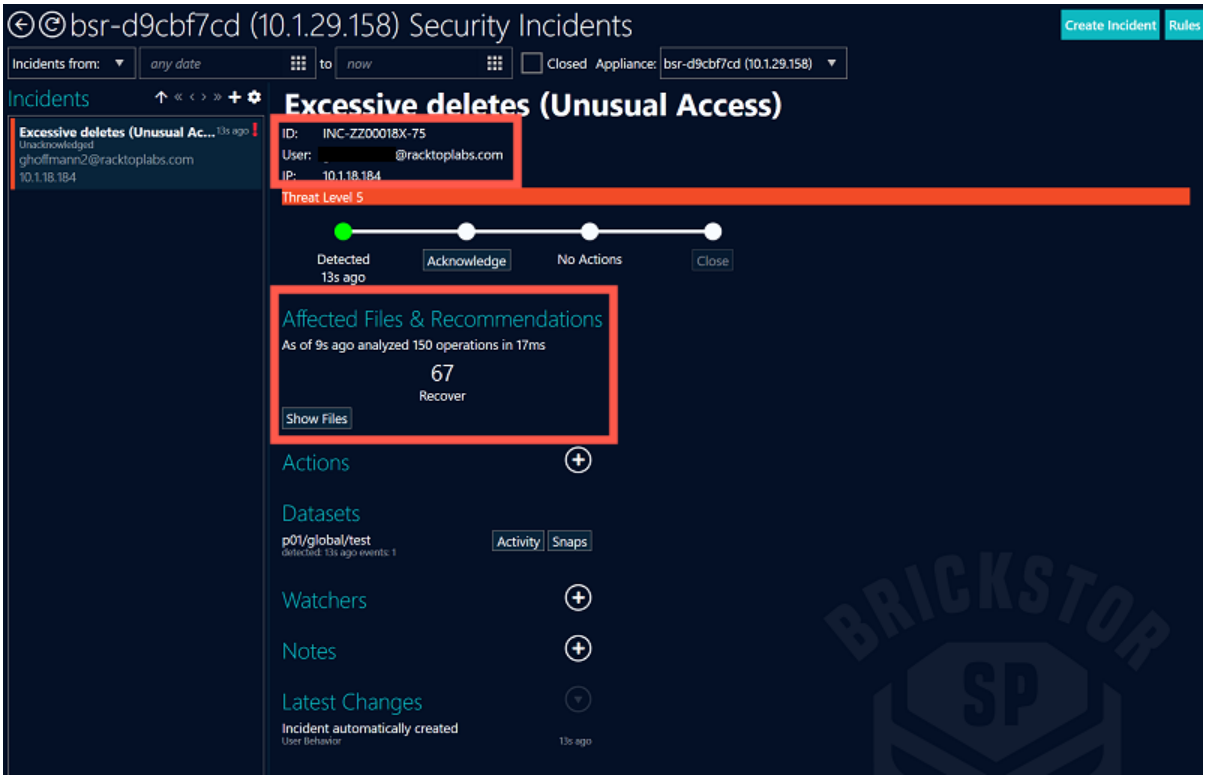

After the file threshold for **Block After** has been reached, the **block** and **hold snapshot** action have been applied for that specific incident.

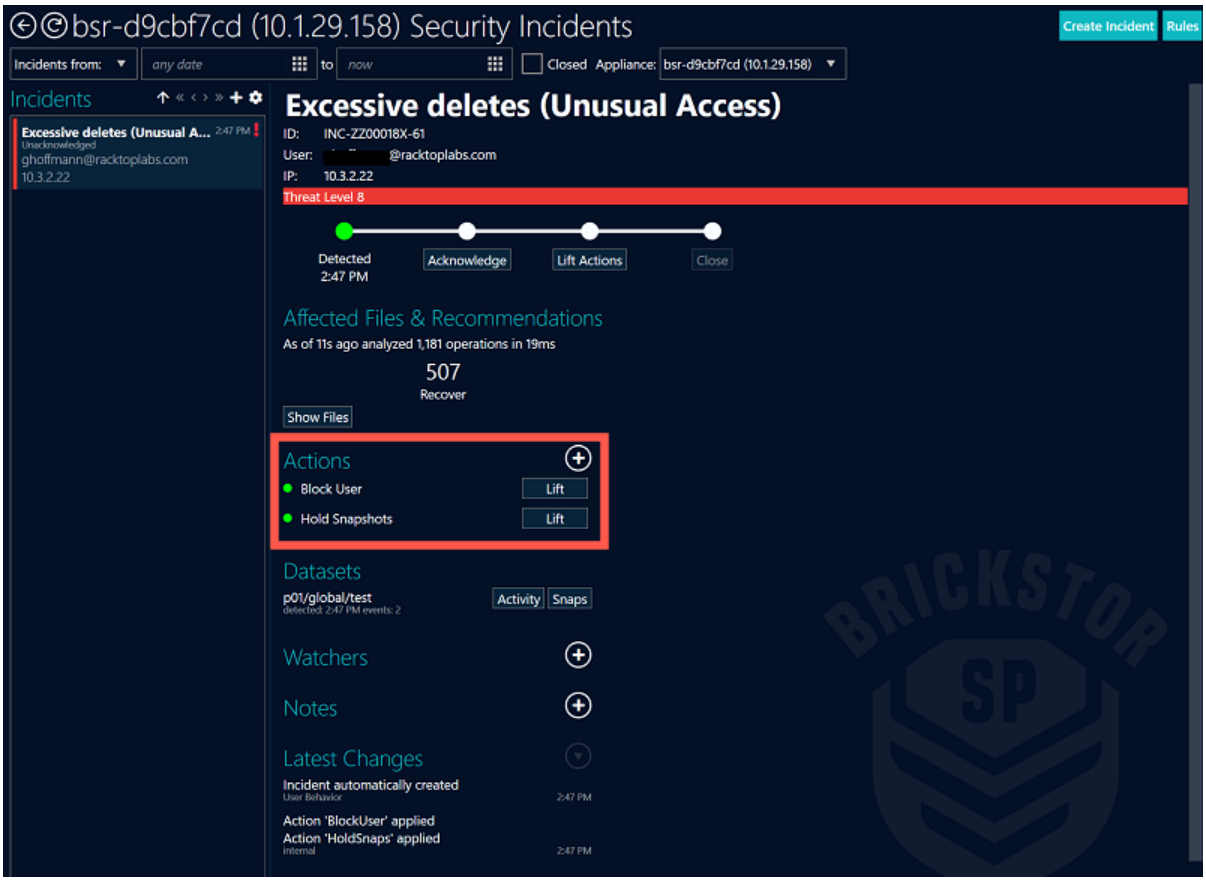

Auto-Reapply

The auto-reapply feature allows the the lifting of any or all actions after an incident has occurred. This will prevent those actions from being reapplied for a specific amount of time.

This allows the performance of normal operations after an incident without being blocked out of the share for a certain amount of time. This will authorize normal actions to be taken shortly after an incident without being flagged.

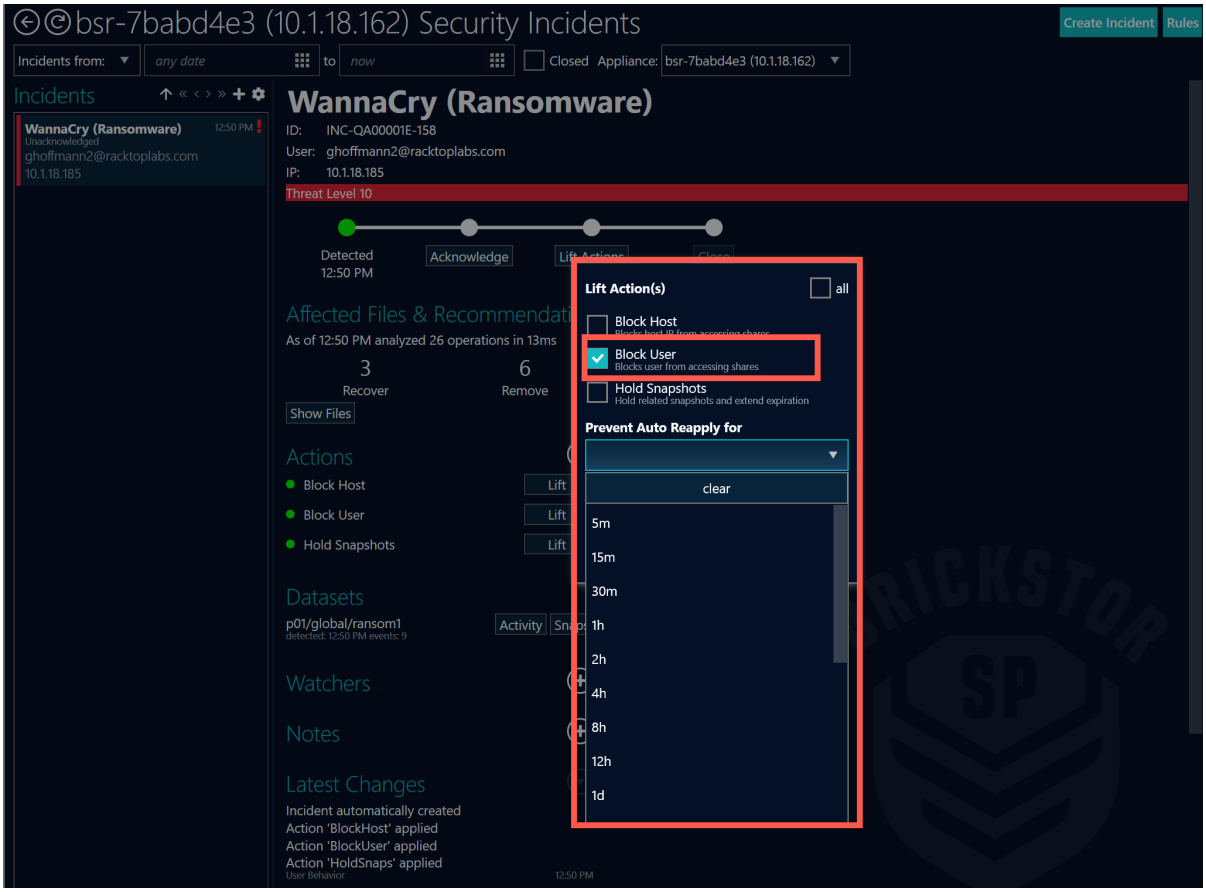

Excessive File Access Assessors

The excessive file access assessors can detect various file operations that stand out in quanitity over a given timespan given typical access patterns. The actions in these assessents include file read, write, and delete operations. Enabling any of the excessive file access assessors can be done in the sharing tab on each dataset. These assessors are configurable on a per-dataset basis.

Enabling the Excessive File Access option will open a new dialog box that allows you to configure how many file operations you want to track per minute. For each of the three file operation trackers, there are options for **Notify after** and **Block after**.

Once the **Notify after** threshold for a certain file operation has been reached, an incident will be created which can be viewed on the **Security Incidents** screen. From here you can see the type of incident, the user and the host IP from which the activity originated and the dataset that was affected.

After the file threshold for the **Block after** has been reached you will see the block and hold snapshot actions have been applied for that specific incident.

Administrator Access Assessors

Administrator access assessors detect when any administrator, domain administrator, enterprise administrator or account operator initiates an operation against a file. The rules for the Admin Access incidents are default rules and the only action applied will be the **Hold Snapshots** action when this incident is triggered.

Once an Admin Access incident is triggered you will see the user account name and the IP address of the device they were using at the top of the screen. You will also see the affected dataset(s) listed, as well as the number of affected files and the **Show Files** option to recover any files if necessary.

Admin Access Incidents

The Active Defense feature also includes an Administrator Access assessor that will detect when any Administrator, Domain Admin, Enterprise Admin or Account Operator does an operation against a file.

The rules for the Admin Access incidents are default rules and the only action applied will be the **Hold Snapshots** action when this incident is triggered.

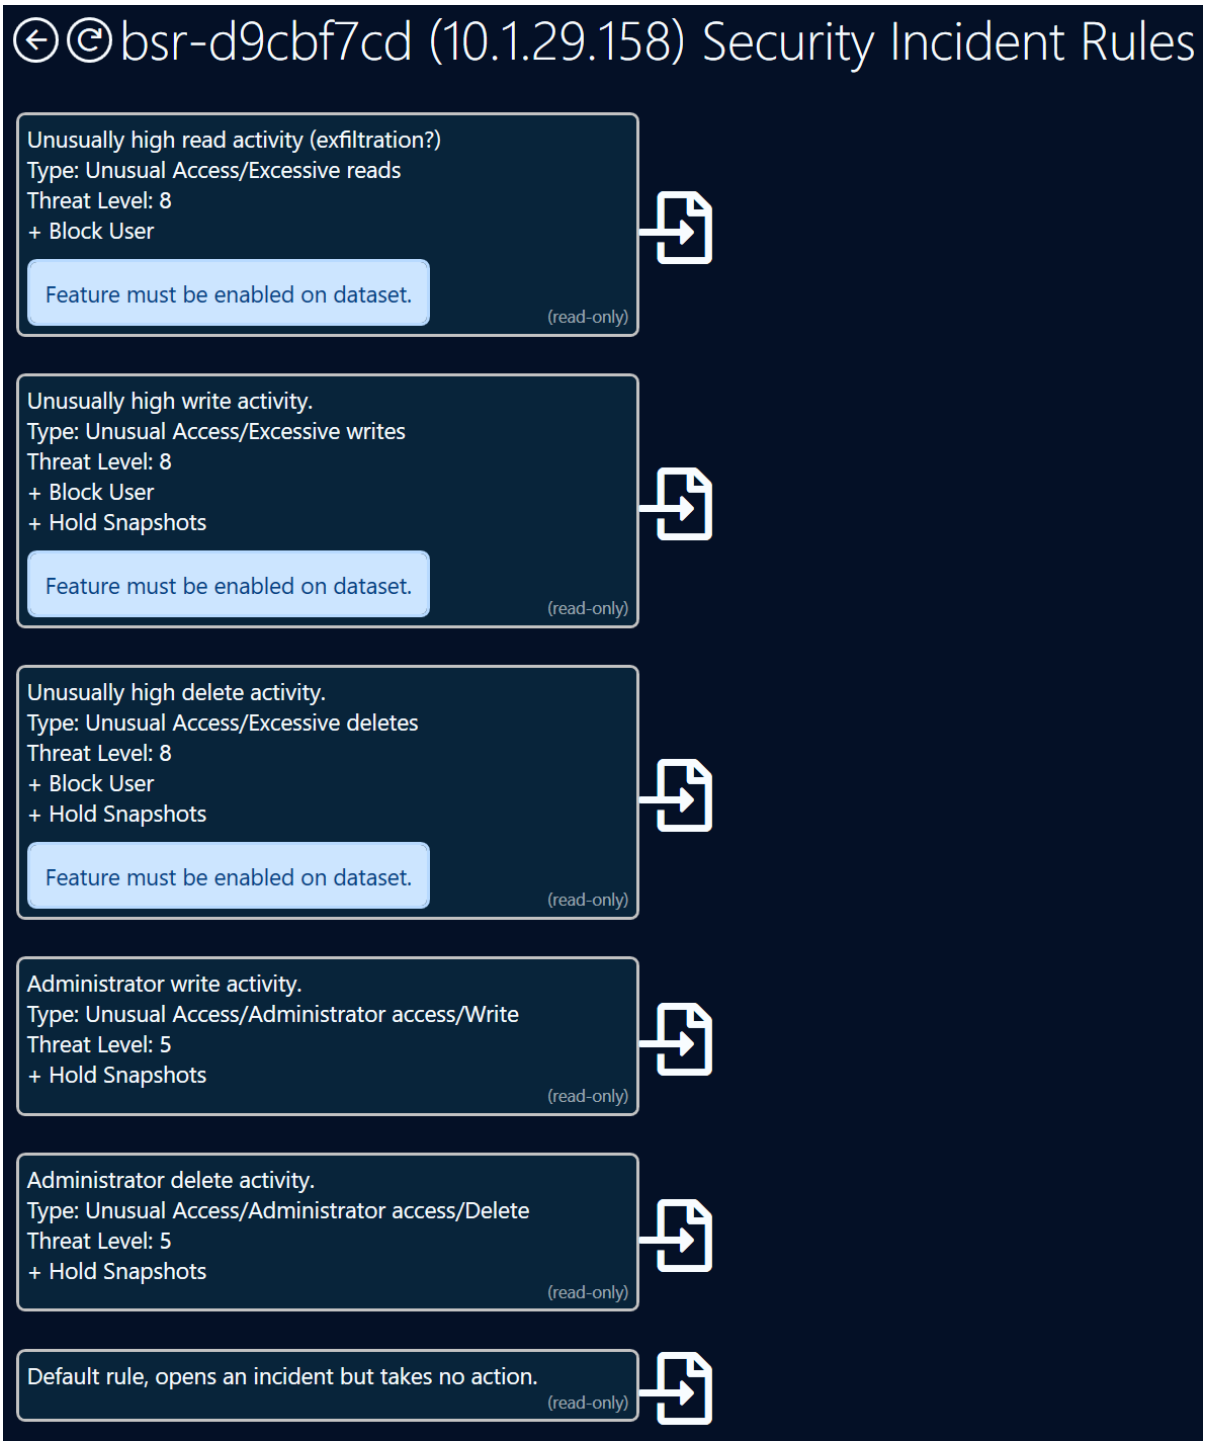

Once an Admin Access incident is triggered, the user account name and the IP address of the device they were using at the top of the screen. The system will also display the affected dataset(s), as well as the number of affected files.

Select the **Show Files** option to recover any files if necessary.

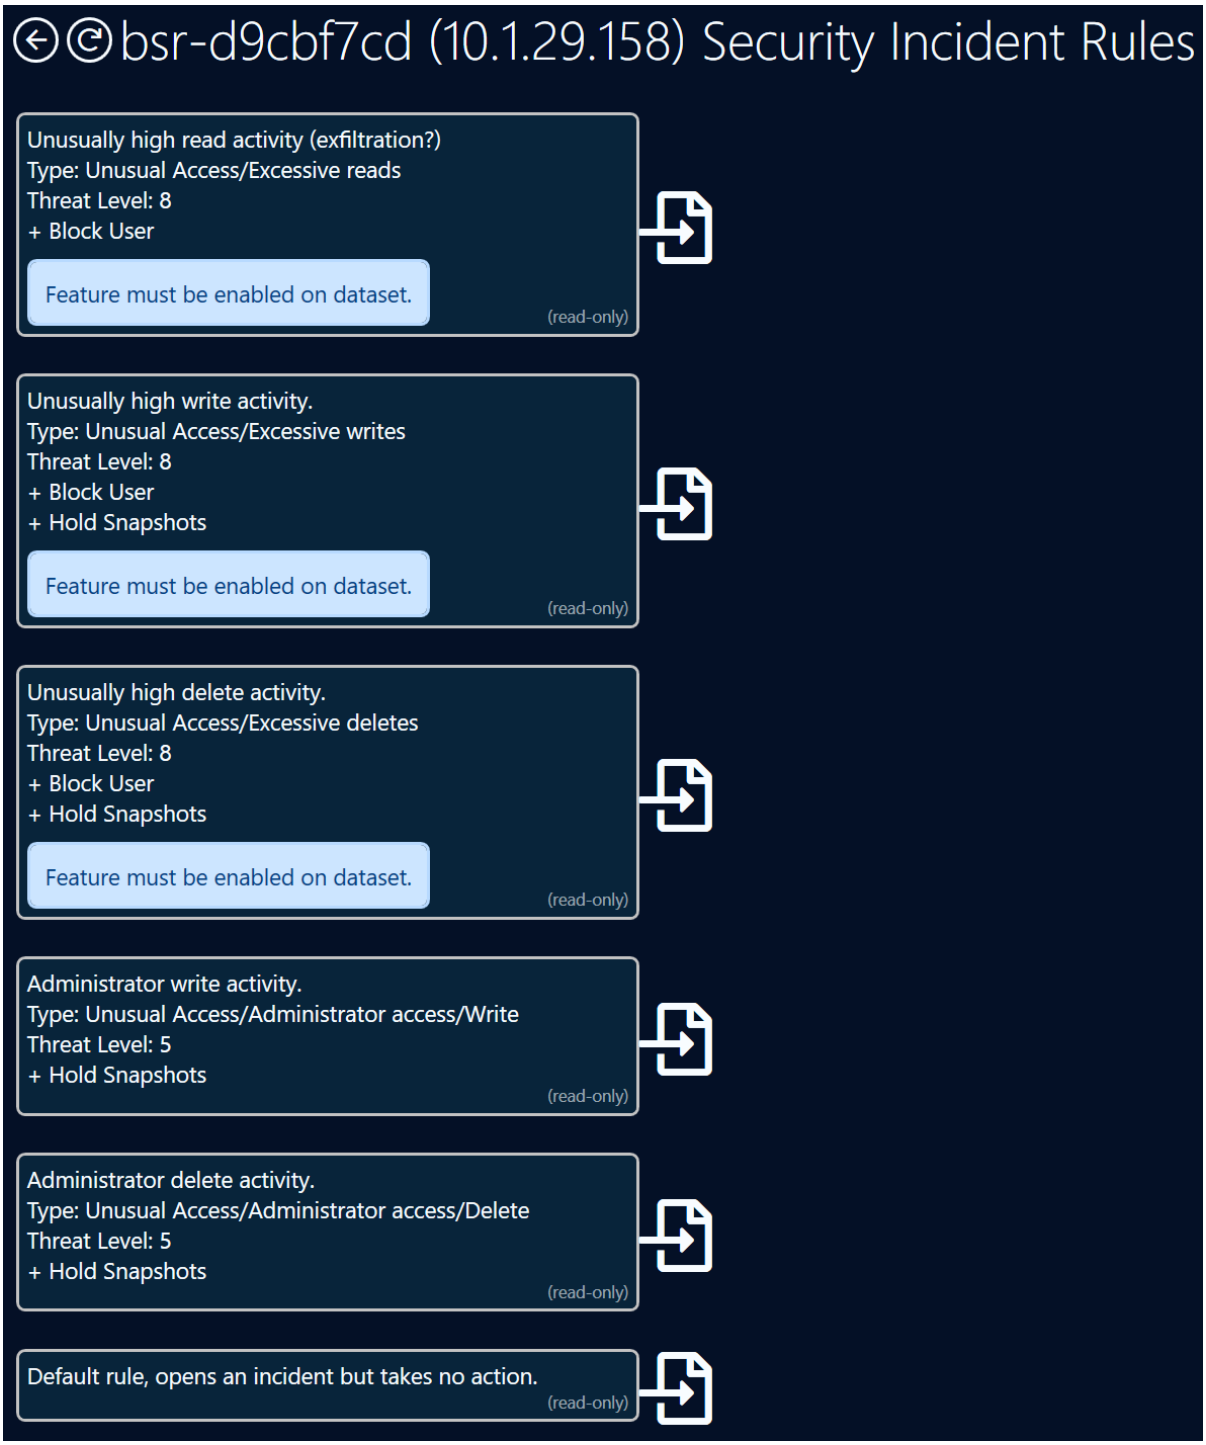

Threat Level

When an incident occurs, threat level is listed in the **events** section in the **incidents** tab. * Threat level indicates the attack's severity via a numbered severity scale (0-10, with 10 being a critical threat), multiplied by the system's confidence in the attack's validity (0-1.0, with 1.0 being absolutely confident).

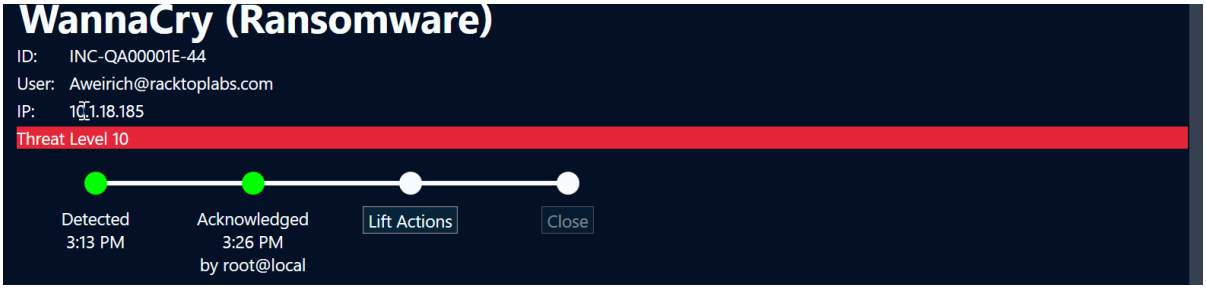

File Recovery

After an incident occurs, the administrator has the ability to see which files have been affected and can decide which ones should be recovered or deleted.

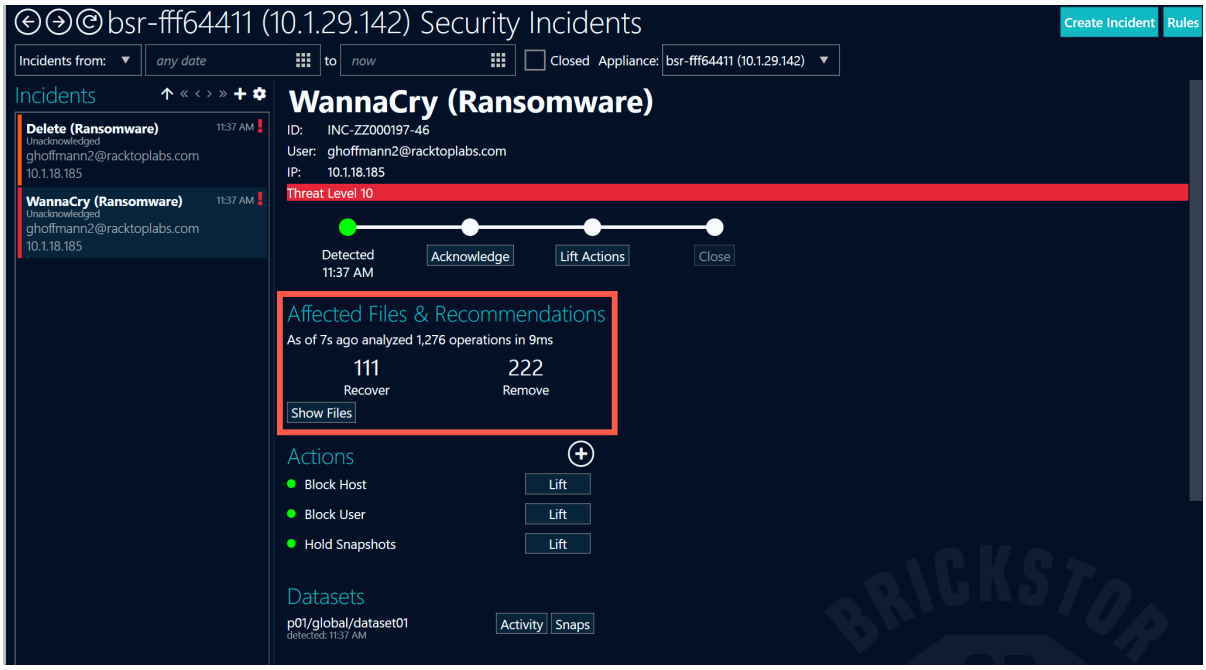

- Click **show files** to see which files have been affected.

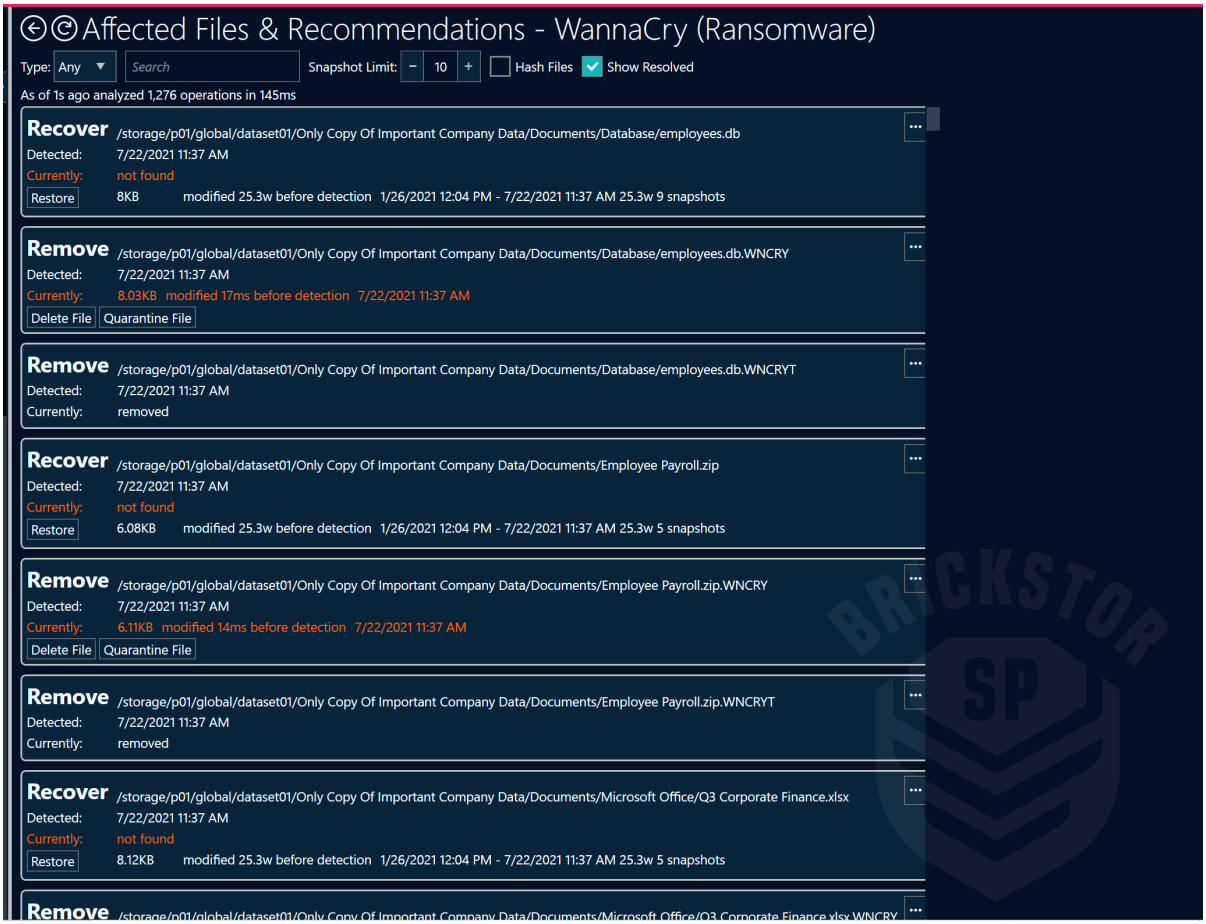

- When accessing the affected share, the administrator is also able to see which files have been affected and the ransom note if one has been added.

NOTE In this case, each file has a WNCRY extension added to it.

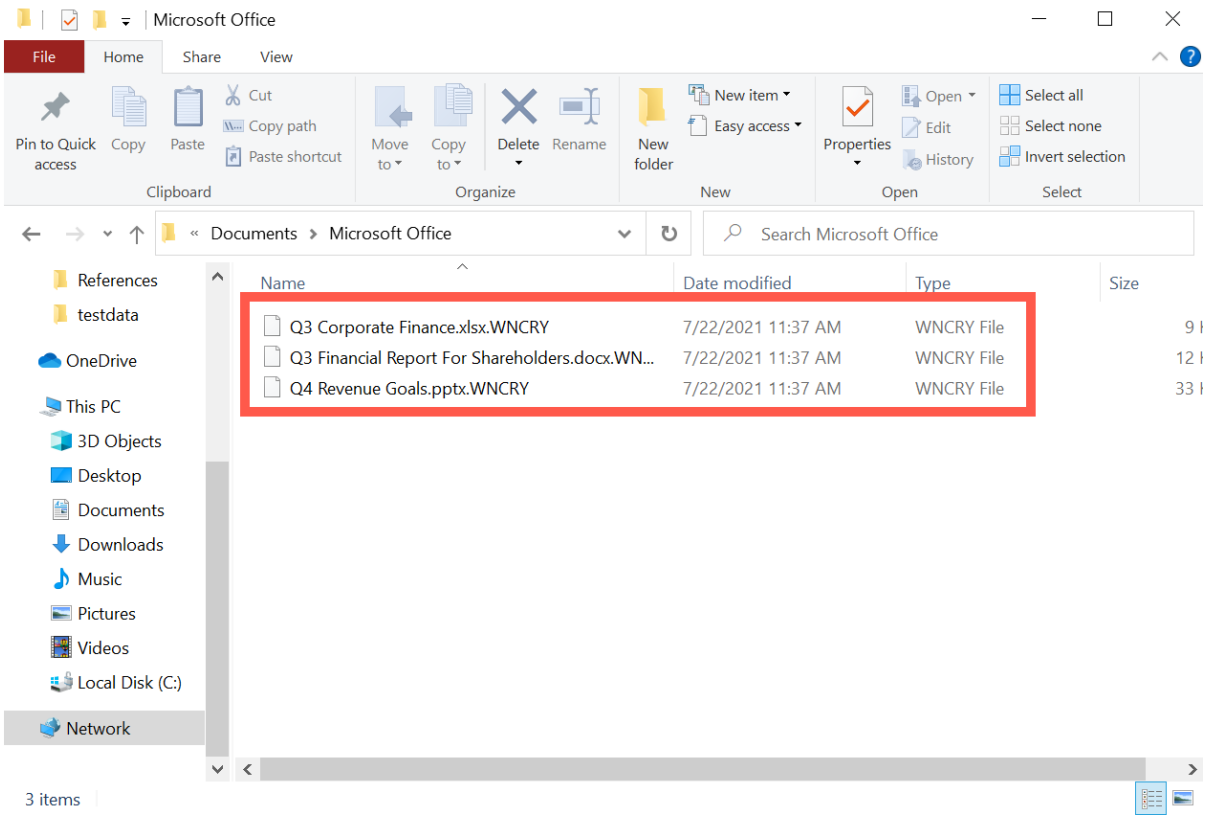

- Click **Restore File** to recover the original file before it was encrypted and click **Delete File** to delete the encrypted file.

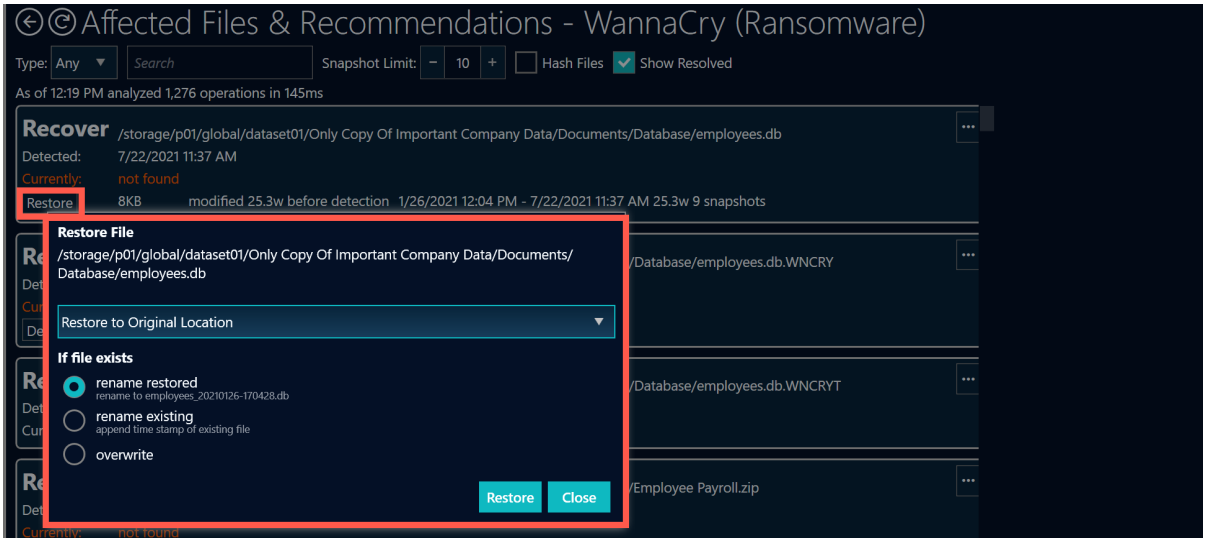

- If a file already exists with the same file name within the share, then there is an option to overwrite that file, rename it to the existing time stamp, or rename it with the current time stamp.
- After restoring the original file and deleting the encrypted file, the share should now only have the recovered version.

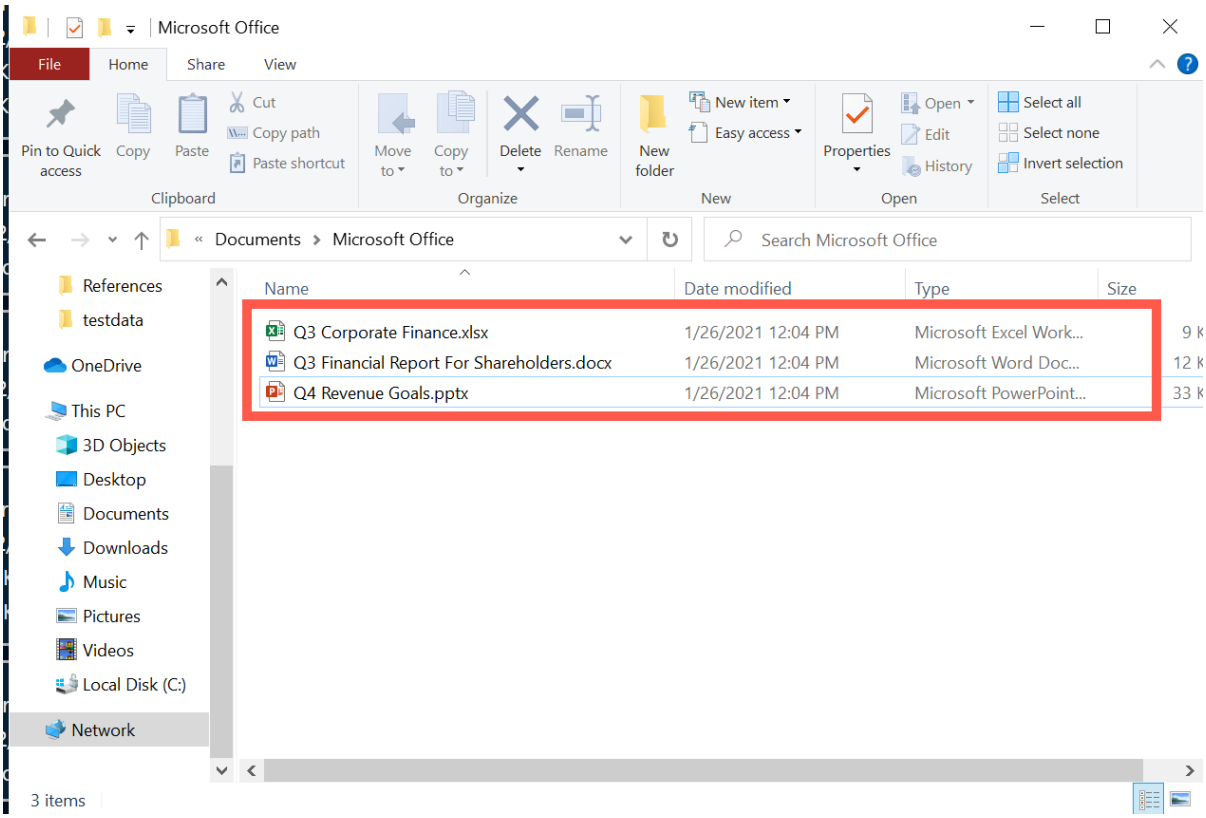

Quarantining a File

There is also an option to quarantine a file in case a file was unable to be recovered correctly.

To quarantine a file:

- Click **quarantine** and it will go into the quarantine dataset located under the **global** dataset.
- Once in the quarantine dataset, go into the **Sharing** tab and enable **SMB share**.

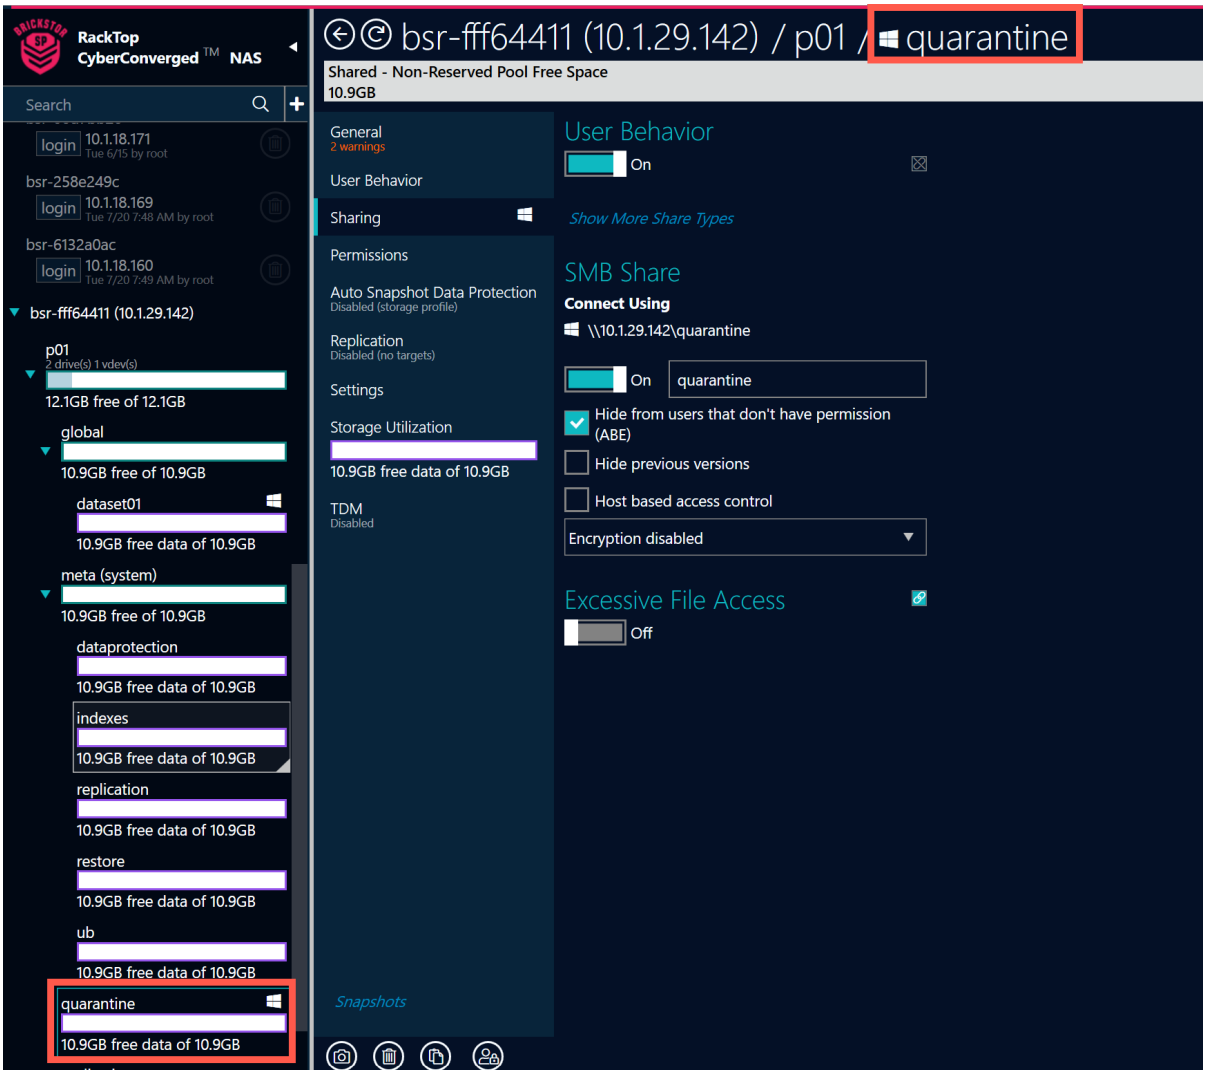

- Hover over the **Connect Using** and click **Go** to access that share.
- Within that share, the file the administrator has chosen to recover can be accessed.
  - This allows the administrator to inspect files before deciding to delete them.

Bulk Recovery

Once an incident has been created, the BrickStor SP Manager will display a breakdown of affected files. This is located by clicking the incident, and navigating to the **Impacted Files and Recommendations** section.

In this section, all affected files will be counted. To see the affected files in further detail, click **View All**.

Immediately below, the affected files are sectioned by whether files were added or removed. If the incident involved unrecognized access, click the **Other** section to view affected files.

Click **View All** to view the files that were impacted by the incident, the following screen will present:

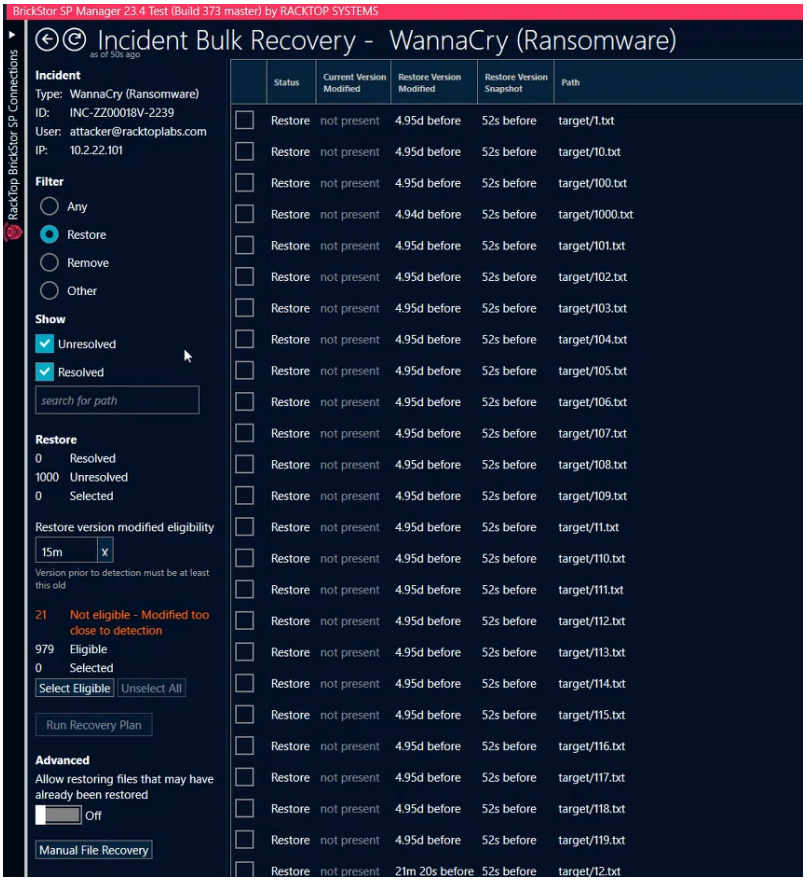

A list of incident-affected files will be displayed. The left-side panel provides options to select/filter through files and control the parameters of the bulk recovery feature.

Bulk Recovery Visual Feedback

Files before and after the recovery process are color-coded to show the progress of the recovery process.

At the Incident Screen, the file breakdown will show a file count of recovered/un-recovered files:

- Recovered files will be displayed in **green** text.
- Removed files will be displayed in **red** text.
- Unresolved files will be displayed in **white** text.

At the Bulk Recovery screen:

- Files that have been **recovered** will be shown with a **green** background.
- Files that are **unresolved** will show with a **dark blue** background.
- Files that have been **removed** will show with a **red** background.

Filters

The filters allow for management of affected files when deciding to recover or remove them based on the type of incident created:

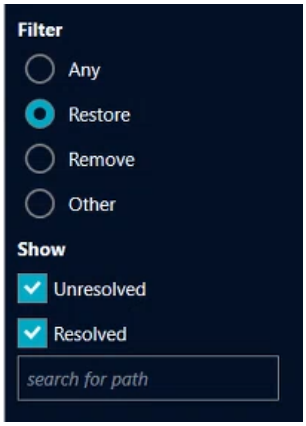

The following options may be selected:

- Any - List all impacted files.
- Restore - List only files with restore recommendation.
- Remove - List only files with remove recommendation.
- Other - List only impacted files without a recovery action (ex. files from an Admin Read incident).

Restore Version Modified Eligibility

The **Restore Version Modified Eligibility** textbox configures the required confidence that the snapshot just prior to the incident is the best snapshot to recovery from, and thus is a good candidate for bulk recovery.

Confidence is configured by adjusting the duration between when the previous file version (before the incident) was last modified and when the file was first impacted by the incident. By increasing the configured duration, a higher confidence requirement is set that the file version just prior to the incident is the best version to recover from.

For example, there is a high confidence that the file version just prior to the incident is the best

recovery version for a file that had not been modified for an hour, up until the file was affected by the incident.

On the other hand, there is a lower confidence that a file version just prior to the incident is the best version to recover from for a file that was regularly modified every second prior to the incident. In the case where a file was modified in close proximity to being affected by the incident, manually choosing the snapshot to recover from is the best option.

Within the supplied text box, a default time of fifteen minutes will be entered.

This value dictates that any file that was edited **within** fifteen minutes of an incident will be ineligible for bulk recovery, and must be recovered or removed manually.

Beneath the modifiable eligibility, information regarding the total amount of eligible and not eligible are displayed numerically.

The buttons below provide an administrator the ability Select/Unselect Eligible files when creating a recovery plan.

Run Recovery Plan

To select files for recovery, manually click checkboxes for the chosen files, or click **Select Eligible** to select all files matching the filters.

Click **Run Recovery Plan** once desired files are selected for recovery.

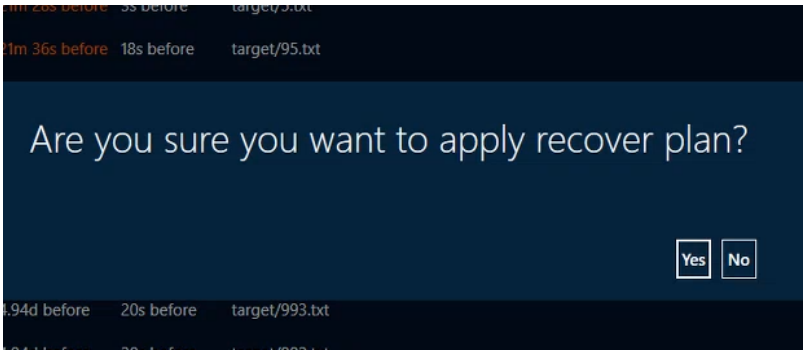

A prompt will present ensuring that the Recovery Plan has been run intentionally:

- Click **Yes** to run the Recovery Plan.
- Click **No** to cancel the Recovery Plan and return to the previous screen.

Clicking **Yes** will initiate the Bulk Recovery of the eligible files. A progress wheel and a numeric description of the recovery progress will present.

Manual File Recovery

Files that prior to the incident were last modified more recently than the configured **Restore Version Modified Eligibility** will not be eligible for the current bulk recovery plan.

To recover those ineligible files, the 'restore version modified eligibility' can be lowered in the next bulk recovery plan, or manual recovery can be used to manually choose the desired recovery version.

To see a detailed view of the ineligible files, and to initiate recovery or removal of the affected files, click **Manual File Recovery**

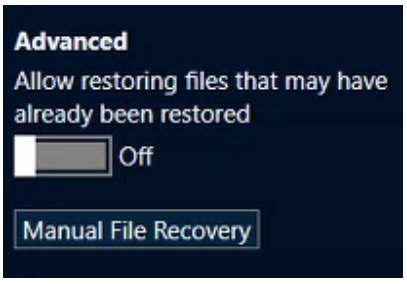

The following screen shows a list of ineligible files, as well as a description of their data and the date/time of the file's last edit.

To augment the list of files, a selection of filters is provided at the top of the screen.

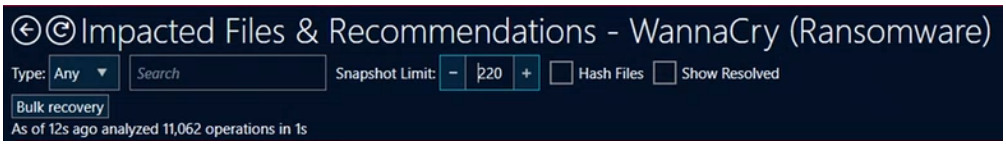

The options provided are as follows:

- Type - Change the list to display any affected file type.
- Search - Search for a specific file.
- Snapshot Limit - Augment the amount of snapshots shown within the list.
- Hash files - Select to show/hide Hash Files
- Show Resolved - Select to show/hide resolved files.
- Bulk Recovery - Select to return to the Bulk Recovery screen.

To initiate a recovery of an affected ineligible file, navigate to the desired file within the list provided.

NOTE A selection of file versions will be listed within each respective file's description. This shows the date/time that the file was edited for each version, as well as a counter for when the file was edited in relation to the incident's creation.

After deciding which file version to select, click **Restore** to initiate the Manual Recovery of the file.

A prompt will display that directs where the file is to be recovered to. Select the appropriate option, then click **Restore**.

NOTE If a file is recovered, but remains encrypted or damaged it is possible to recover an older version by following the steps above, and selecting an older version of the file. The option to **Overwrite** should be selected to replace the file with a version that is usable.

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

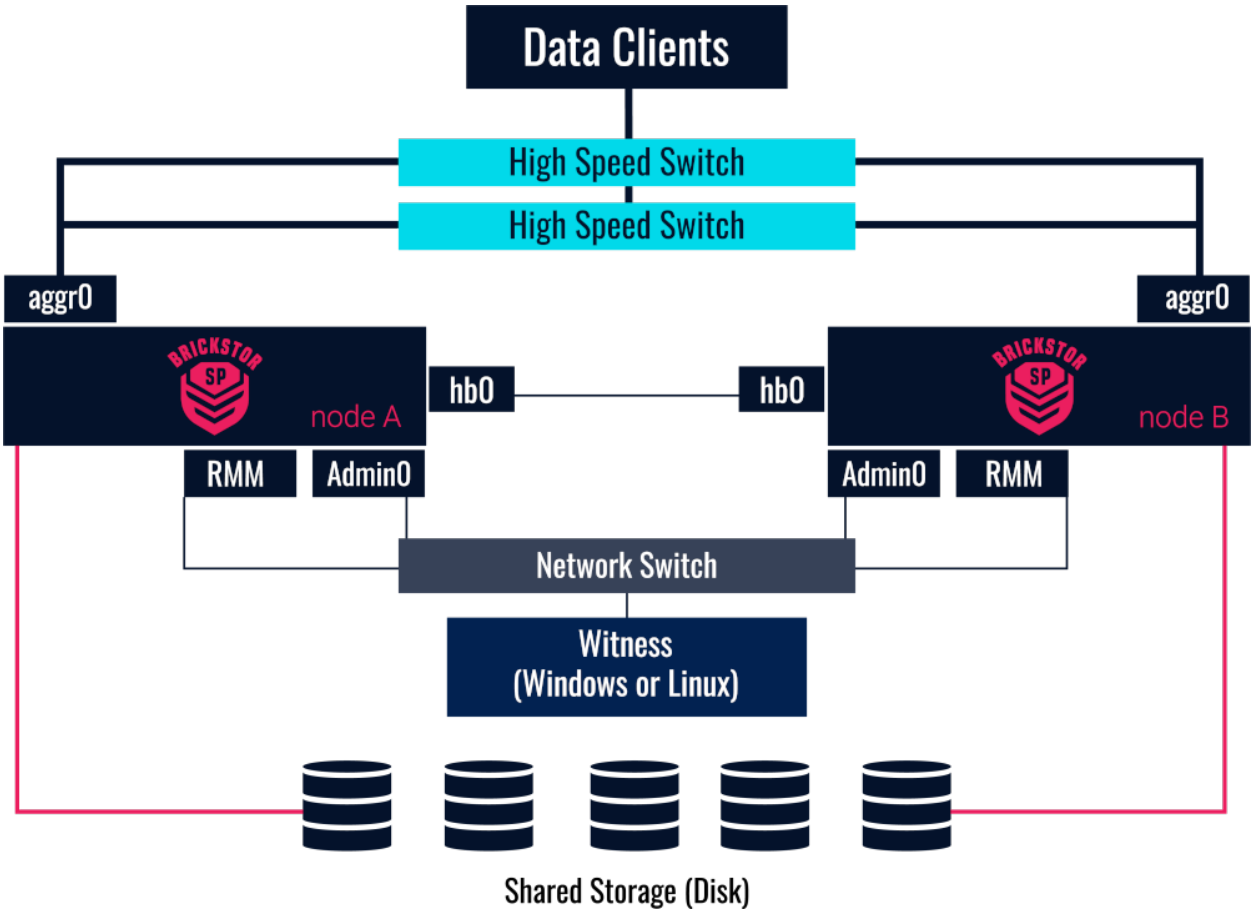

HA Scenarios

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 exibiting 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 further instruction into the cluster setup wizard, the following prerequisites must be met in order to form a BrickStor SP 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.

Configuring Heartbeat

BrickStor SP HA uses a heartbeat interface directly connecting both nodes over a dedicated network interface. In order to configure heartbeat interface do the following:

Establishing Connection

- Choose one Ethernet port on each BrickStor SP node.
- Connect those ports together using a standard network patch cable. Modern servers do not require crossover cables.

NOTE	When the network connection is established, some hardware may not activate the port LED until the interface is configured in the OS.

Creating Heartbeat Interface

- Using BrickStor SP CLI, list available physical interfaces by running: (Option `-m` displays MAC address which can be helpful to identify server's network interfaces)

```
dladm show-phys -m
LINK         SLOT     ADDRESS            INUSE CLIENT
igb0         primary  0:1e:67:80:1d:8c   yes  --
igb1         primary  0:1e:67:80:1d:8d   no   --
igb2         primary  0:1e:67:80:1d:8f   yes  --
igb3         primary  0:1e:67:80:1d:8f   no   --
```

- Configure heartbeat over the desired interface. For example, to use "igb2" interface run:

```
dladm create-vnic -l igb2 hb0
```

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>:

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>:
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) Dependencies

The following prerequisite installations must be satisfied to install Witness on a Linux workstation:

- `bzip`
- `ipmitool`
- `dmidecode`

1. To install them, enter the following command into your terminal:

```
sudo yum install bzip2 ipmitool dmidecode -y
```

2. Configure witness system's firewall to allow inbound TCP port 4746 to communicate with BrickStor SP nodes. The below example uses firewalld.

```
firewall-cmd --list-all
firewall-cmd --permanent --zone=public --add-port=4746/tcp
firewall-cmd --reload
firewall-cmd --list-all
```

NOTE	In cases where outbound traffic is blocked, allow ports tcp/4746 and tcp+udp/623 from Witness to both BrickStor SP nodes.

Install Witness (Linux)

WARNING	Ensure that installation prerequisites are satisfied (see above section).

Download the Witness

- Open your BrickStor's IP address using a web browser.
  - url: https://<brickstor_ip>:

BRICKSTOR
SP
CyberConverged™ NAS
BrickStor SP 22.0.2.201
Use of this system requires acceptance of the End User License Agreement
Terms, Conditions, and Warranty information
Release Notes and Configuration Guide
Appliance Certificates
You must install and trust both certificates for the click-once application to work.
Download Server Public Certificate
Download Server CA Certificate
High Availability Witness Binaries
Download Witness for Windows
Download Witness for Linux
BrickStor SP Manager Client
brickstorspmgr-22.0.2-19.zip
RACKTOP®

NOTE

Optionally, the Witness package can be downloaded using Secure Copy Protocol.

```
scp root@xx.x.xx.xxx:/usr/racktop/bsrapid/static/witness/ha-witness-linux-xx.x.x.xxx.tar.bz2 .
The authenticity of host 'xx.x.xx.xxx (xx.x.xx.xxx)' can't be established.
RSA key fingerprint is SHA256:H3cUrhhTTTRSP1WBWSbh8gXRb8CrZmoIzIUFaNyaTgI.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'xx.x.xx.xxx' (RSA) to the list of known hosts.
Password:
ha-witness-linux-xx.x.xx.xxx.tar.bz2
```

Decompress the Witness

```
tar -xf ha-witness-linux-xx.x.xx.xxx.tar.bz2
```

Copy the hiavd binary to **/usr/sbin**

```
cp /root/ha-witness-xx.x.xx.xxx/hiavd /usr/sbin
```

Make the binary executable.

```
chmod 555 /usr/sbin/hiavd
```

Configure hiavd as systemd service. Edit the `hiavd.service` file and add the contents below:

```
vi /etc/systemd/system/hiavd.service
```

```
[Unit]
Description = RackTop Systems High Availability Daemon (hiavd)
After = syslog.target network.target
[Service]
Type                = simple
ExecStart           = /usr/sbin/hiavd -w /var/run -pid /var/run/hiavd.pid
ProtectHome         = true
ProtectSystem       = true
Restart             = on-failure
RestartSec          = 2
StandardOutput      = journal
StandardError       = journal
WorkingDirectory    = /var/run
[Install]
WantedBy = multi-user.target
```

Enable hiavd.service to run on boot.

```
sudo systemctl enable hiavd.service
```

Start hiavd for the first time.

```
sudo systemctl start hiavd.service
```

NOTE After creating your cluster using Brickstor SP Manager, the hiavd.service may exit. There will be an error in the gui stating the witness is faulted. If this happens, start the service again.

Service Control

The service can be stopped or restarted using standard systemd commands:

```
sudo systemctl start hiavd.service

sudo systemctl stop hiavd.service

sudo systemctl restart hiavd.service

sudo systemctl status hiavd.service
```

High Availability requires time to be in-sync for all nodes as well as Witness. To enable Network Time Protocol (NTP):

- Install the NTP service.
- Modify the NTP configuration file, **'/etc/ntp.conf'**, with required options.
- Add **reference clock peers** to the configuration file.
- Add **drift file** location to the configuration file .
- Add optional statistics directory to the configuration file .
- Enable and start the NTP service.
- Check operation and synchronization status.

Distributed Configuration Database (confd) Windows Install

The following steps will guide the Distributed Configuration Database (confd) installation on a Windows system. The install handles all of the steps necessary for the confd instance to run as a standalone node (copying to proper directories, firewall, service install, etc).

NOTE This does not actually join the cluster.

Steps to join the cluster must be followed after installation is complete.

- From the target Windows machine, using a web browser, navigate to **https://IP_OF_YOUR_BRICKSTOR**.
- Login.
- Click the **Download Witness** for Windows.

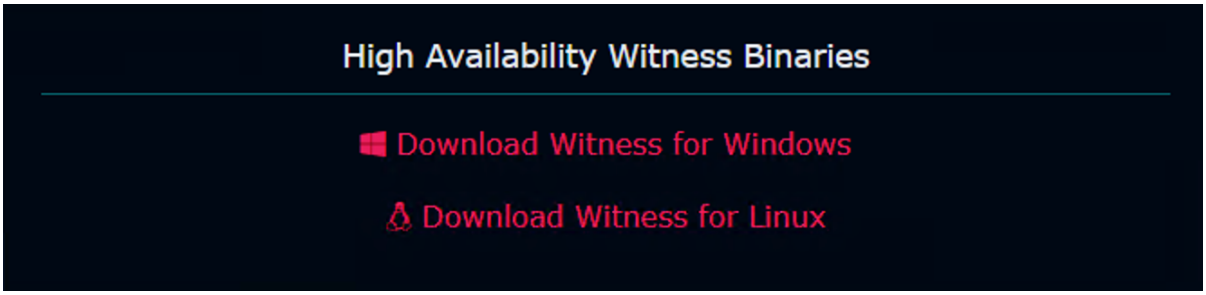

- Unzip the file to the local **Downloads** folder.
- Right-click and choose **Run as Admin**.
- Select **Install** from the menu.

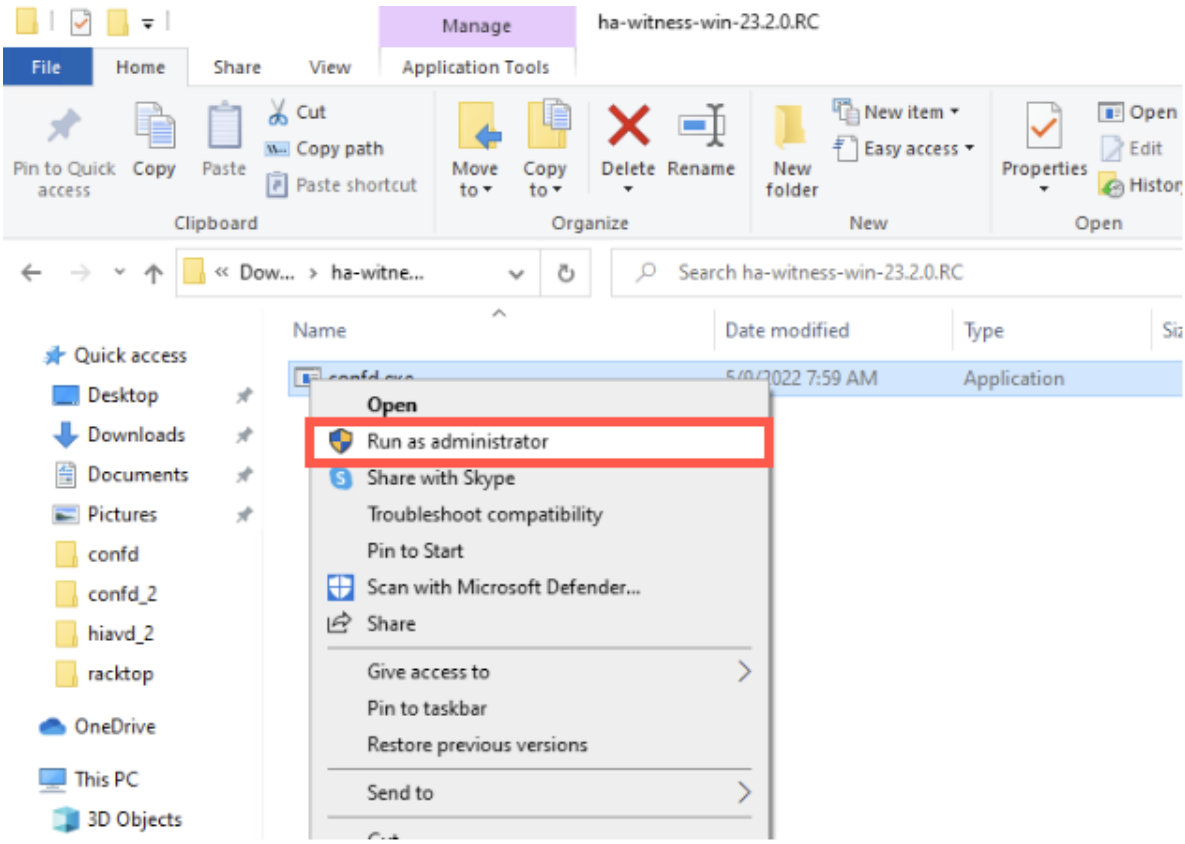

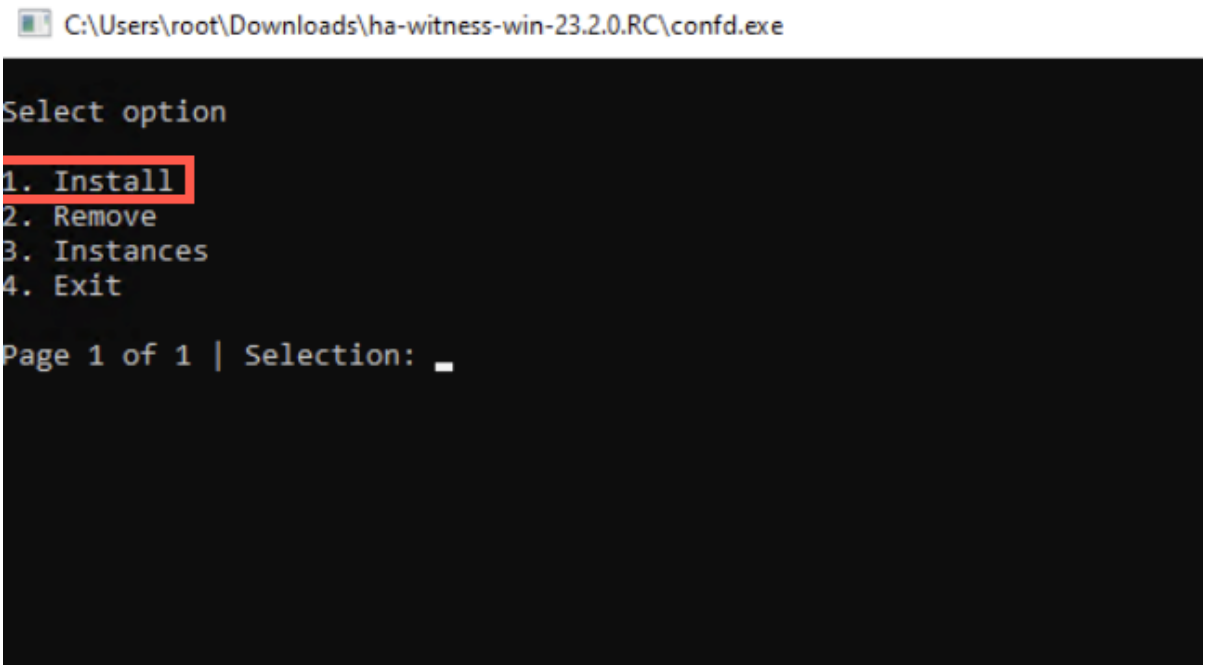

- Follow the prompts to install.

View Existing Instances

- Launch the menu as **administrator**.
- Select **option 3**.

To Remove Existing Instances

- Launch the menu as **administrator**.
- Select **option 2**.

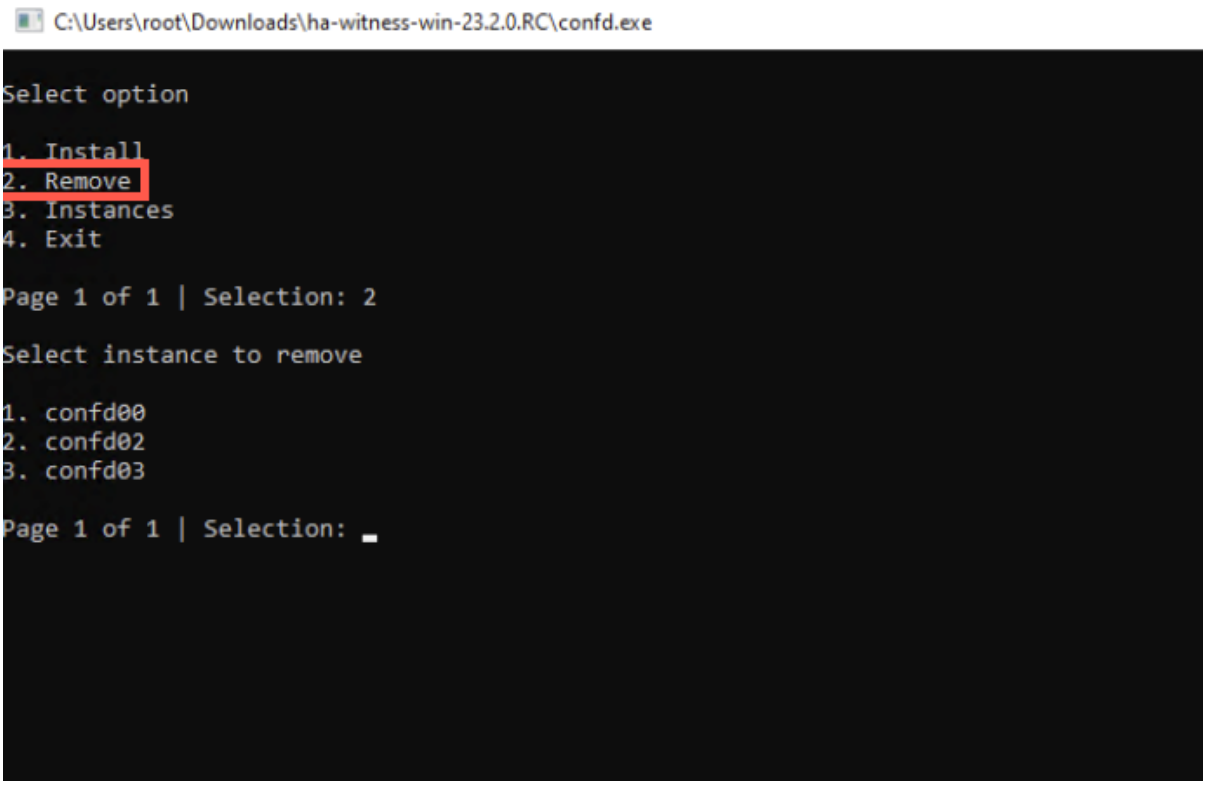

- Select the instance intended for removal.

Setup Confd

- To begin confd setup:

```
confadm -instance <instance number> member show
```

- Ensure the member count is one. A new cluster cannot be joined if this HA cluster or node is already a member of one.

NOTE	The port being used for the PeerUrls.

- Setting peer address number:

```
confadm -instance <instance number> member set-peer-address X.X.X.X:PORT_FROM_STEP_2
```

- To join:

```
confadm -instance <instance number> join
```

NOTE	There is a known bug where a display of setup falure may be shown. Validate true failure by running **confadm member show**. If the return shows more than one member, setup can be continued.

Reset and Start Over

- Run **confd.exe** as admin and choose Remove option to remove the desired instance.
- Delete **C:\Program Files\RackTop\Brickstor\confd\${INSTANCE_NUMBER}** of the instance

just removed.

- Perform the install steps over again.

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

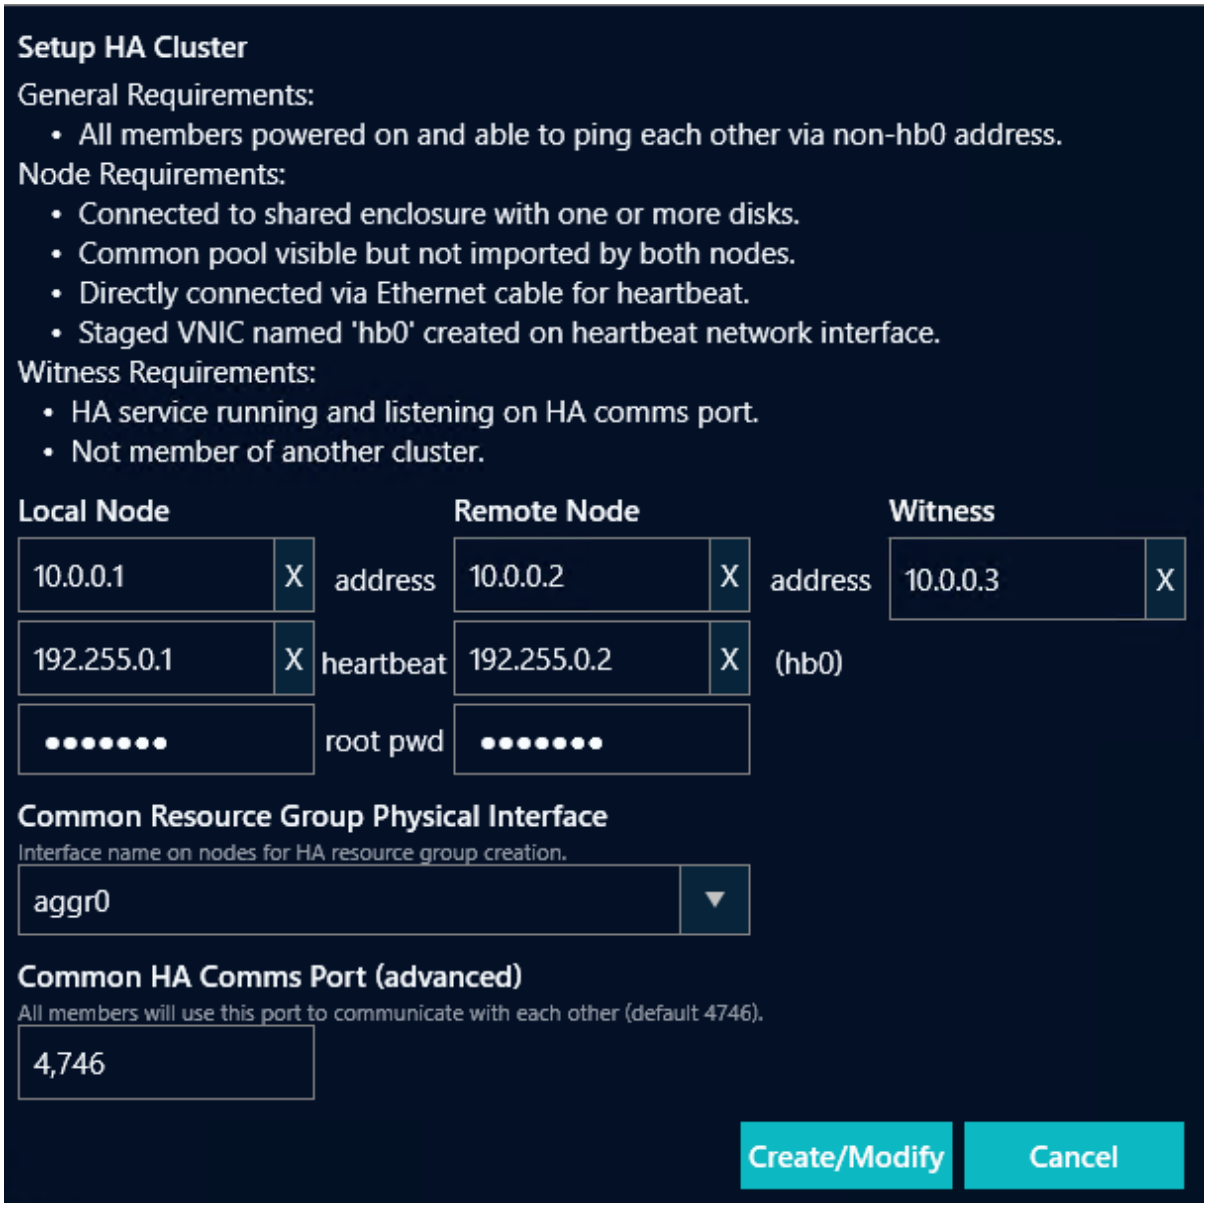

Local Node - the node you are currently managing. Lets call it the first node.

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 3. 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 ballon 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.

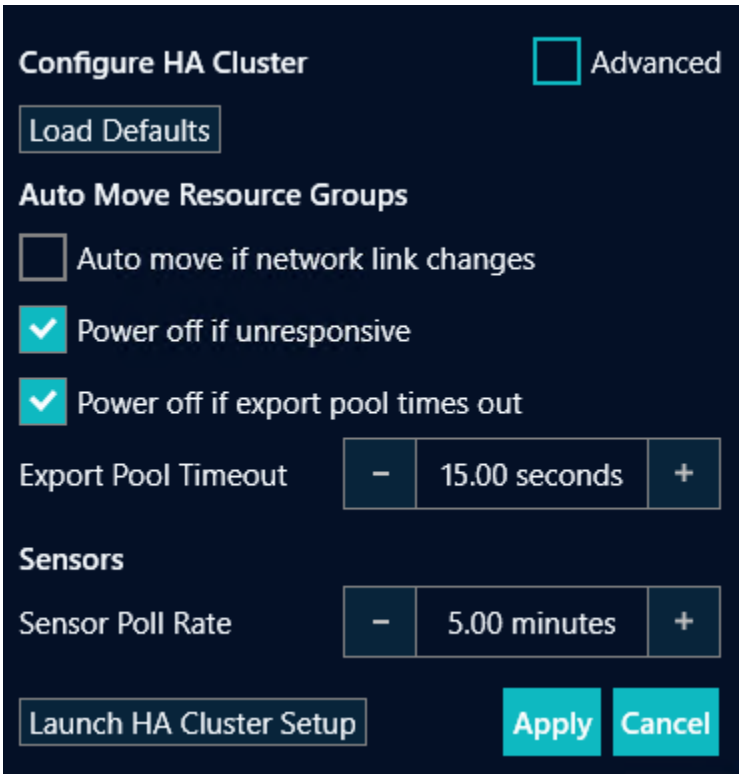

Table 4. HA Settings

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 n 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 n seconds before forcefully powering off failed node. See Power off if export pool times out

Setting	Default Value	Description
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 the **Stop** button . The **Disable HA Node** dialog box will open with additional options.

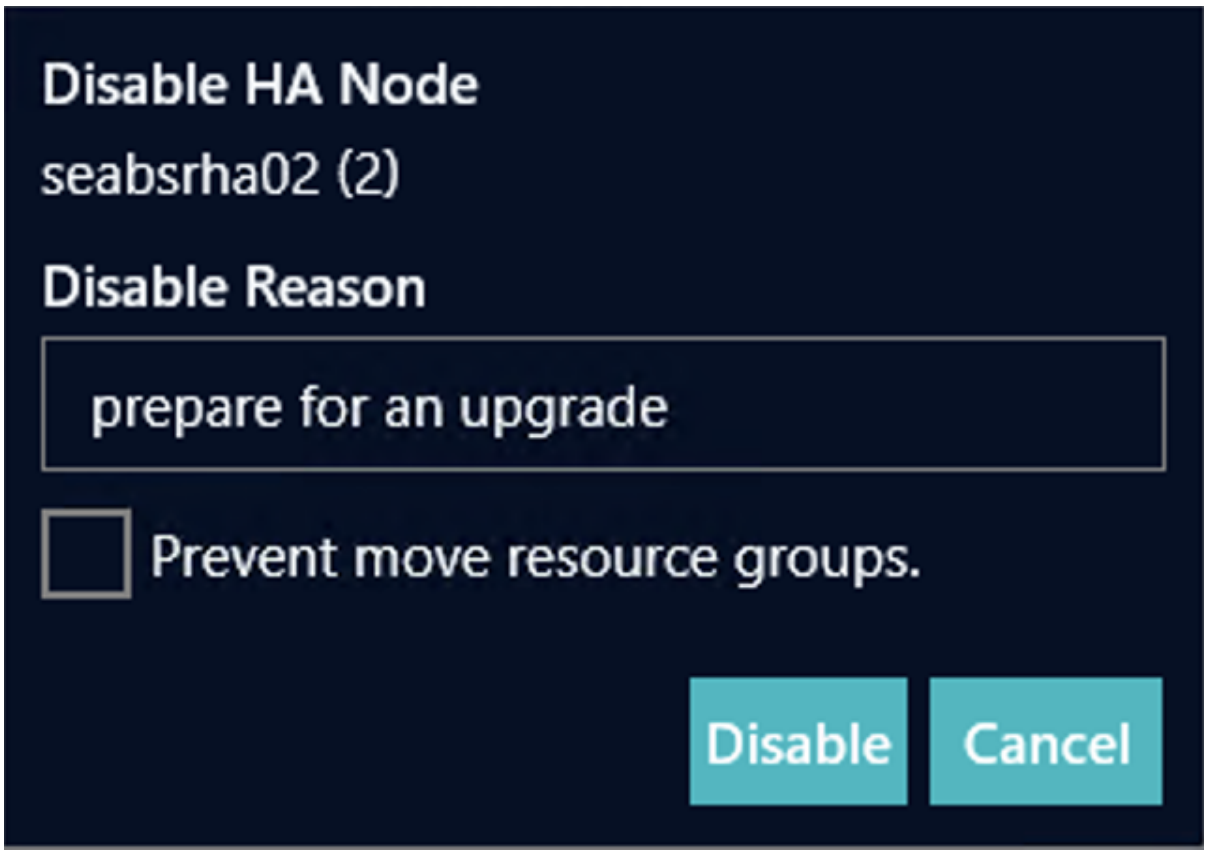

4. In the **Disable HA Node** dialog enter a **Reason** message. This message will be saved in the event log to provide a detailed explanation for this operation.

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.

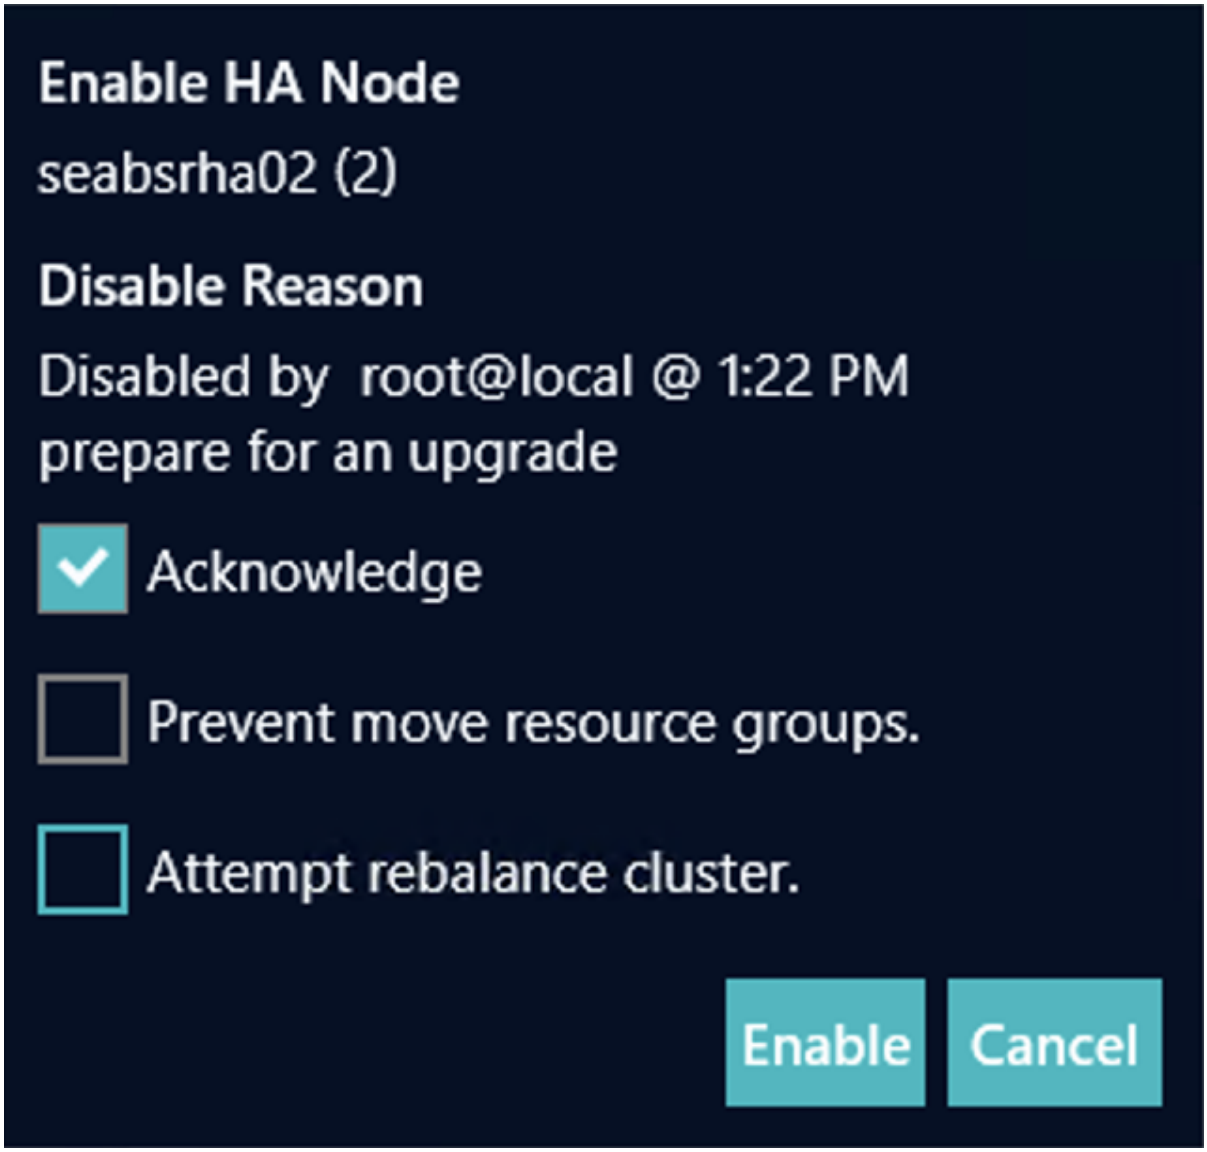

4. In the Enable HA Node dialog check **Acknowledge** box to confirm delivery of the message

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.

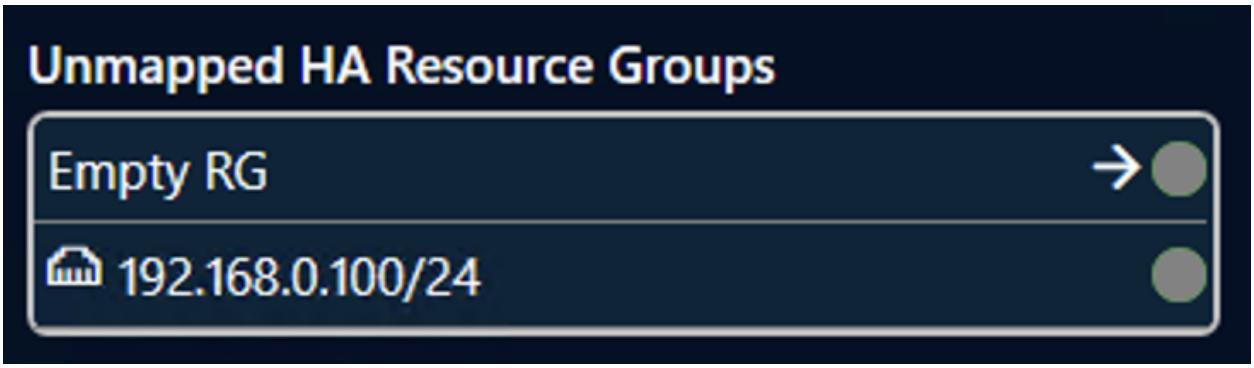

Resource Group States

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.

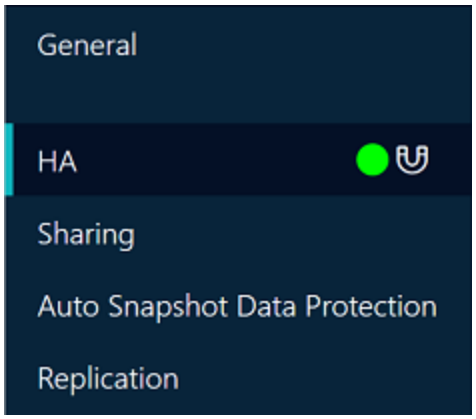

3. Hover over the HA Cluster and then click the plus-R icon to add a Resource Group. This will bring up the HA Resource Group creation dialog.

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:

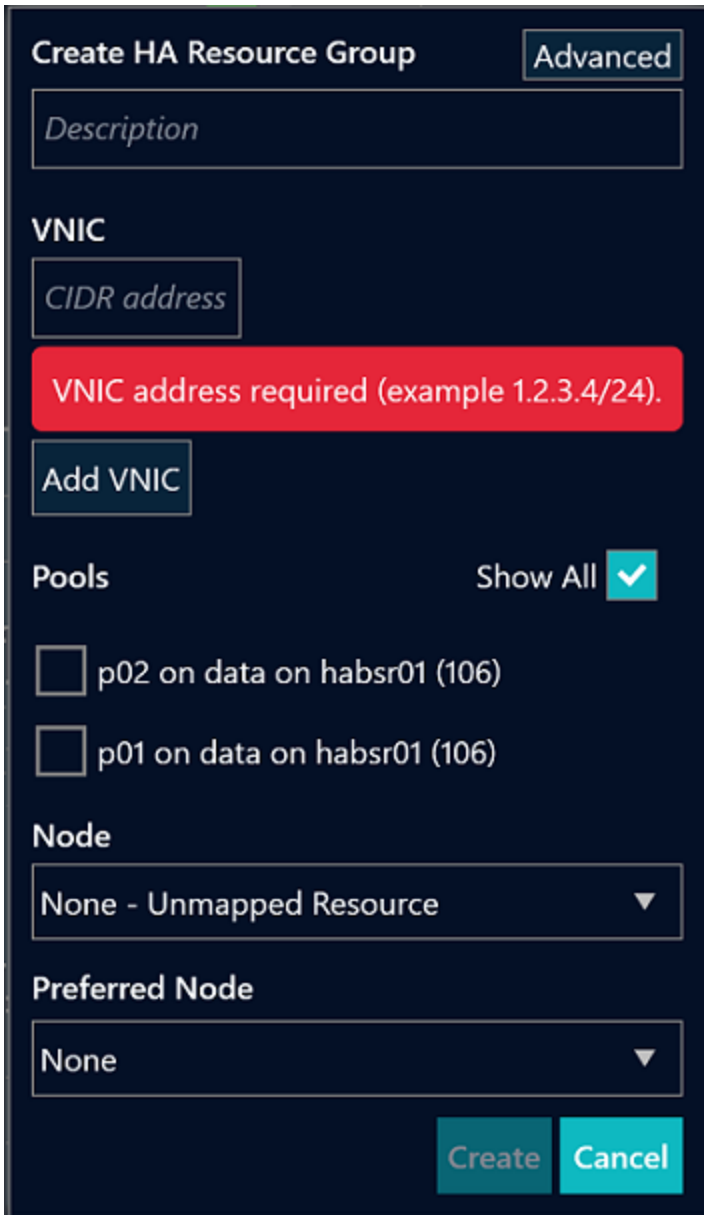

a. **Description** - Enter meaningful text describing the purpose of this Resource Group (ex: "User Data")

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.

5. 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.

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.

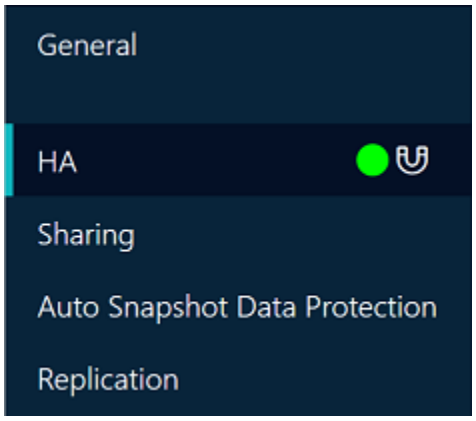

3. Hover over the HA Cluster and then click the plus-R icon to a Add Resource Group. This will bring up the HA Resource Group creation dialog.

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

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

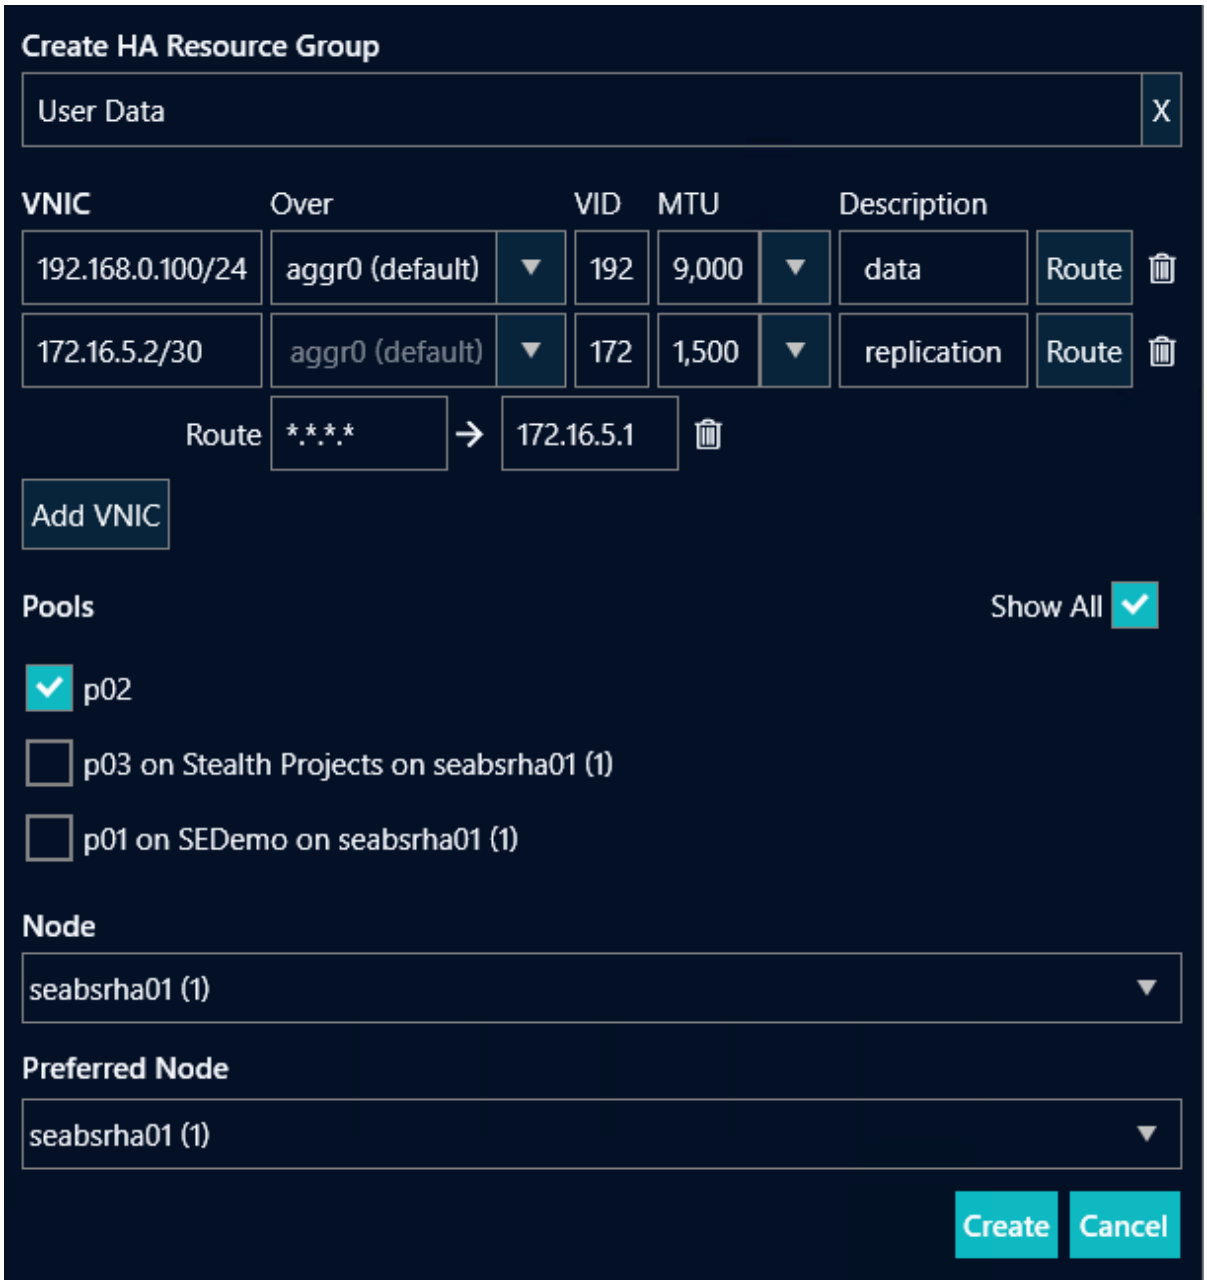

a. Click the **Advanced** button to show advanced property fields.

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

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

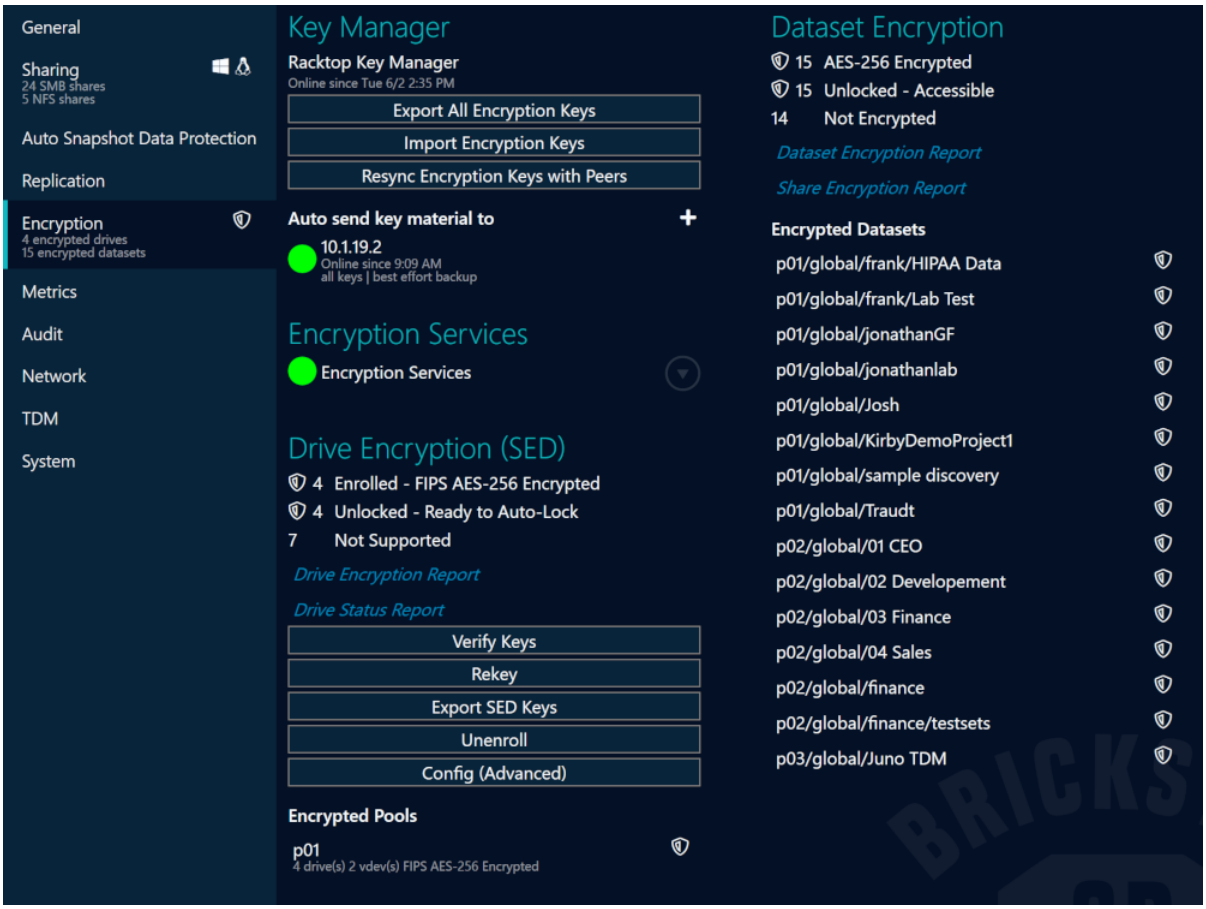

Drive Encryption Related Buttons

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

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.

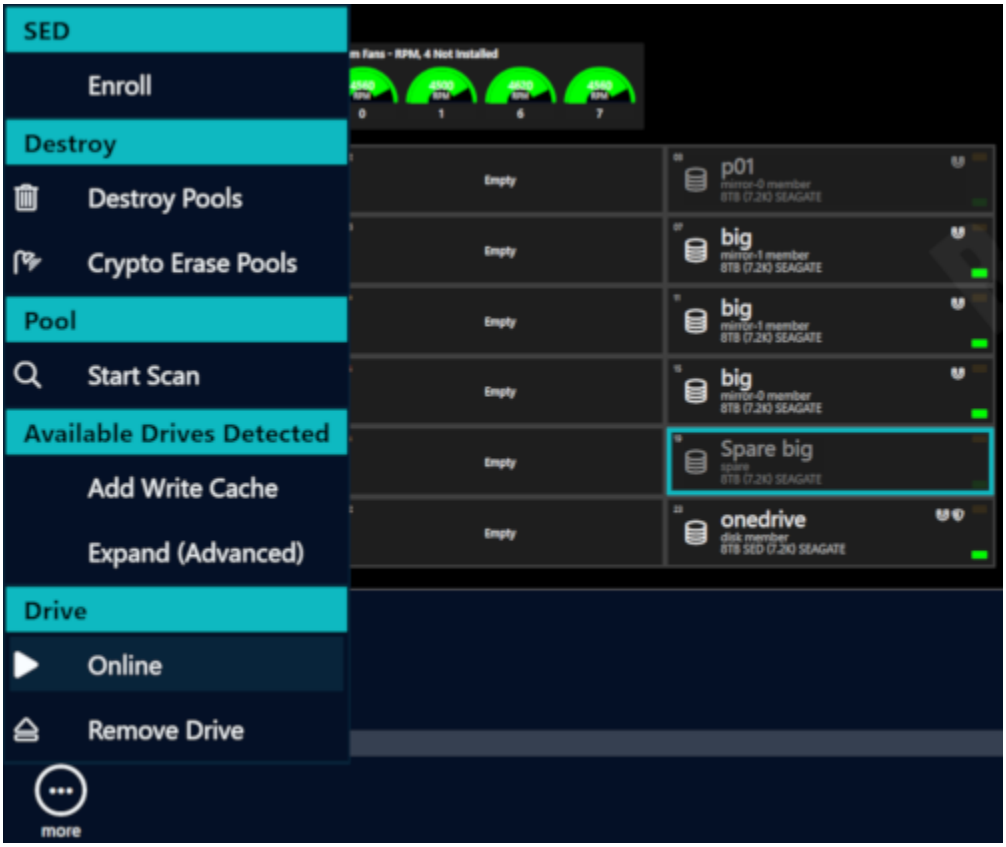

Other Self Encrypting Drive Operations

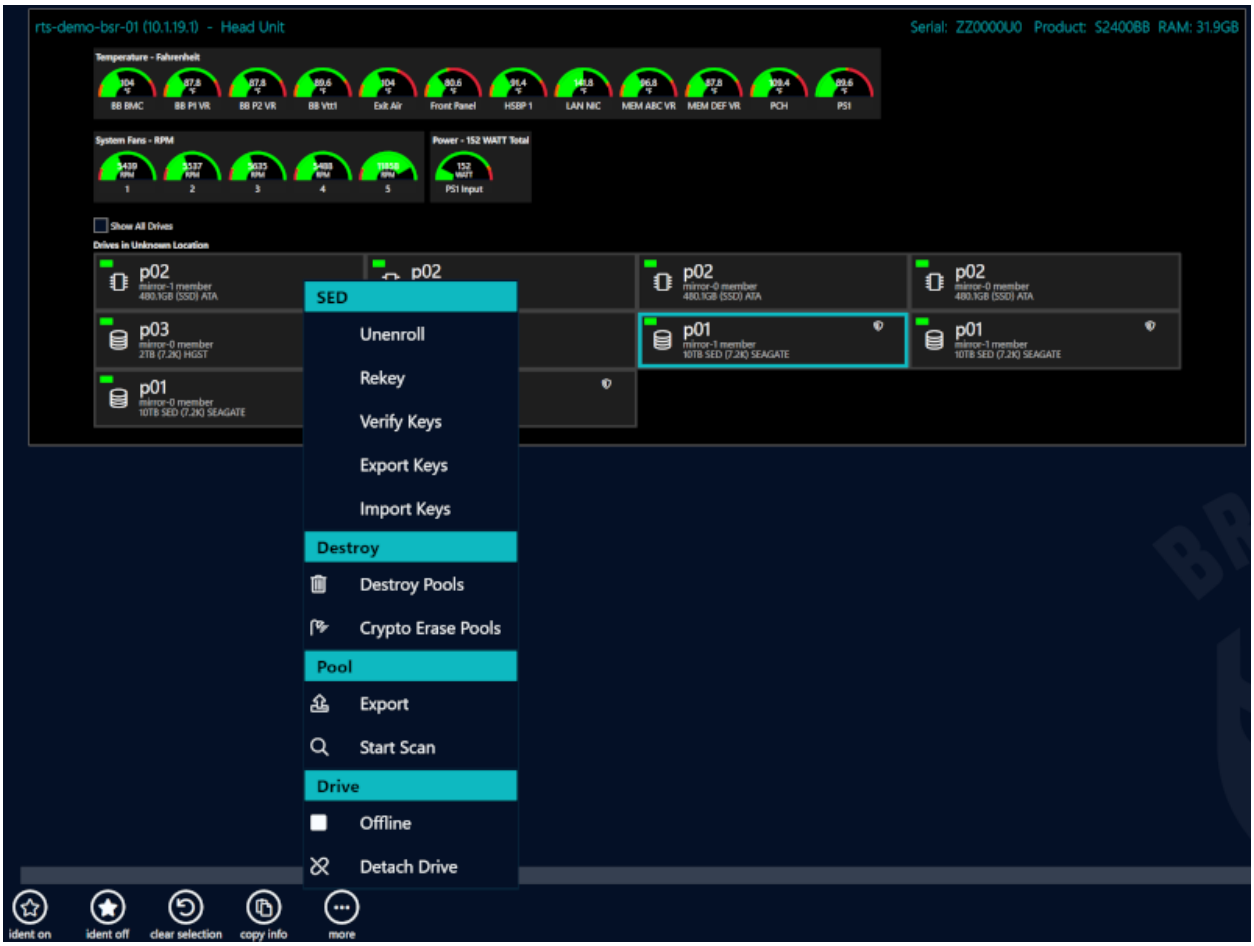

Unenroll – Removes drive from SED management and sets the drive to default PIN and sets the drive to stay unlocked.

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 internal. 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.

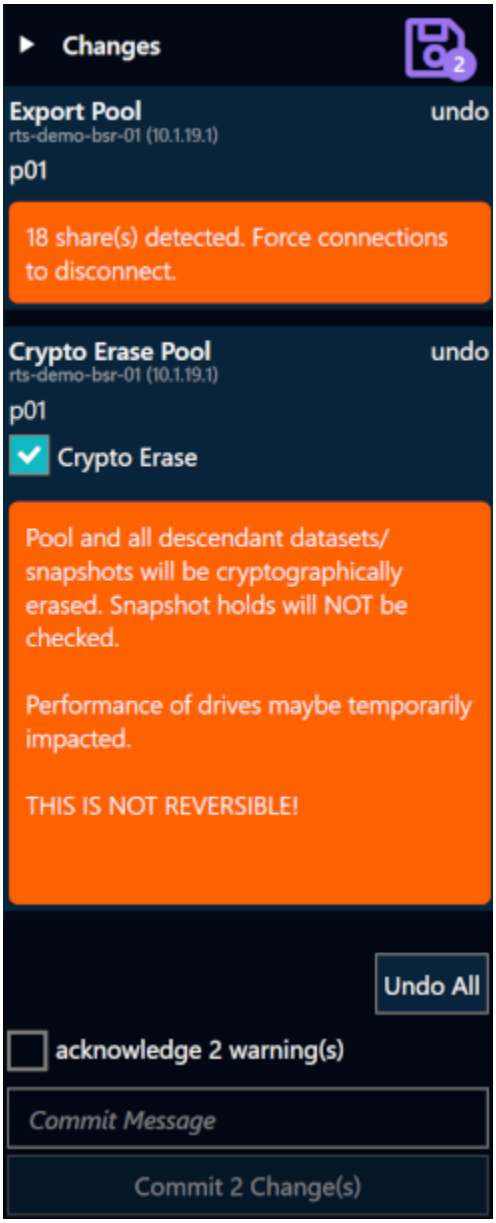

SED Protection on the Main Pane

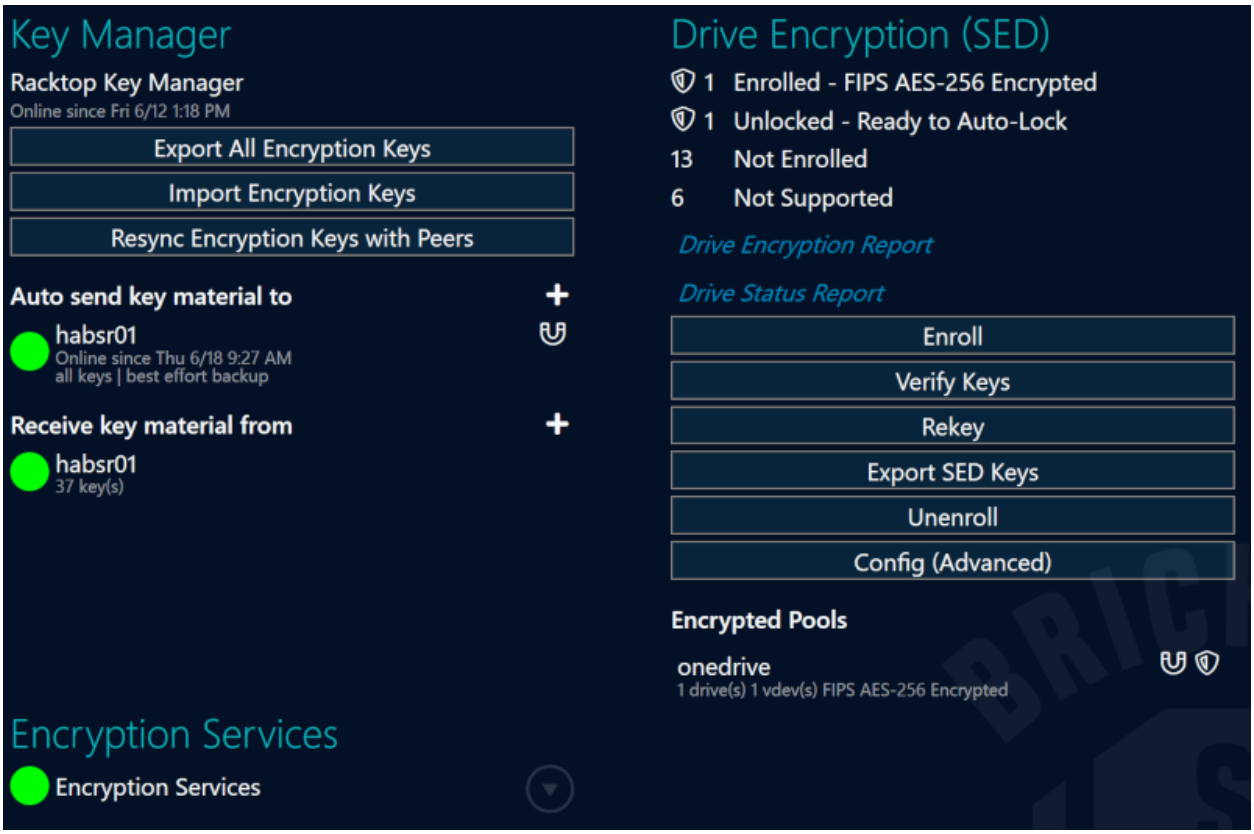

Under the general tab of BrickStor SP Manager users can perform various SED configuration options as well review reports about which drives are enrolled in SED management and the current status of each drive.

Transparent Data Movement (TDM)

Transparent Data Movement (TDM) is a patent pending technology developed by RackTop to enable the seamless movement of data between tiers of storage within a BrickStor SP to external storage tiers including other BrickStor SP nodes, third party NFS capable storage and S3 compliant object storage. TDM is an advanced hierarchical storage management feature of the BrickStor SP operating system that enables policy-based security and compliance to be applied to data stored in the cloud or on third party storage systems. Policies can be applied to the data set to determine which target the data should tier to when policy dictates it should be moved to a more economical tier of storage. Users continue to access data through the same client protocols using the original file path and do not need to change their workflow.

File Chunking

BrickStor SP's TDM feature intelligently chunks large files into smaller objects when tiered to an object store. The benefit of this chunking is that when a large file is updated that has been tiered to an object store, only the chunks with modifications must be updated in the object store. If the file was stored as one large object, then the entire file would have to be retrieved and rewritten as an object. This unique feature of BrickStor SP saves bandwidth and speed of file access and cost when using a cloud-based object store.

Demand Cache

BrickStor SP's intelligent demand cache optimizer reduces cost and improves performance by reducing IO for remote files on the economic tier. When a file is tiered with TDM, it is not actually removed from the primary storage tier, the "demand cache", until the space needs to be reclaimed by the OS for other files. This means if a user opens a file before it is evicted from the demand cache, the file will be opened from the primary tier's demand cache. This eliminates any latency from the economic tier and costs that may be imposed from a cloud provider for IO and data retrieval. The version of the file on the economic tier will be updated if there are any modifications made to the file.

Logical Segmentation – Enclave Elimination

Many organizations want to eliminate physical system segmentation and silos to enable centralized monitoring and dynamic resource allocation. Previous security challenges can be overcome with the advanced access control features built into BrickStor SP's Operating System. BrickStor SP includes granular access control capabilities to restrict access down to the individual file level.

BrickStor SP includes discretionary access control across all client platforms, which are the most common access control scenarios. Discretionary access control is sufficient for government security accreditation of multiple enclaves within the same security domain. BrickStor SP supports host-based access control on top of discretionary access control. With SE Linux and NFS 4.2, BrickStor SP can support mandatory access control through the support of context security labels. With this architecture, a single BrickStor SP system can be accredited for access from multiple security domains and enclaves.

Configuring TDM

To enable TDM on a dataset:

1. Select the dataset in the Connections pane
2. Select the TDM tab in the Detail pane
3. Click the "Enable TDM on Dataset" button

Steps to enable TDM on a dataset

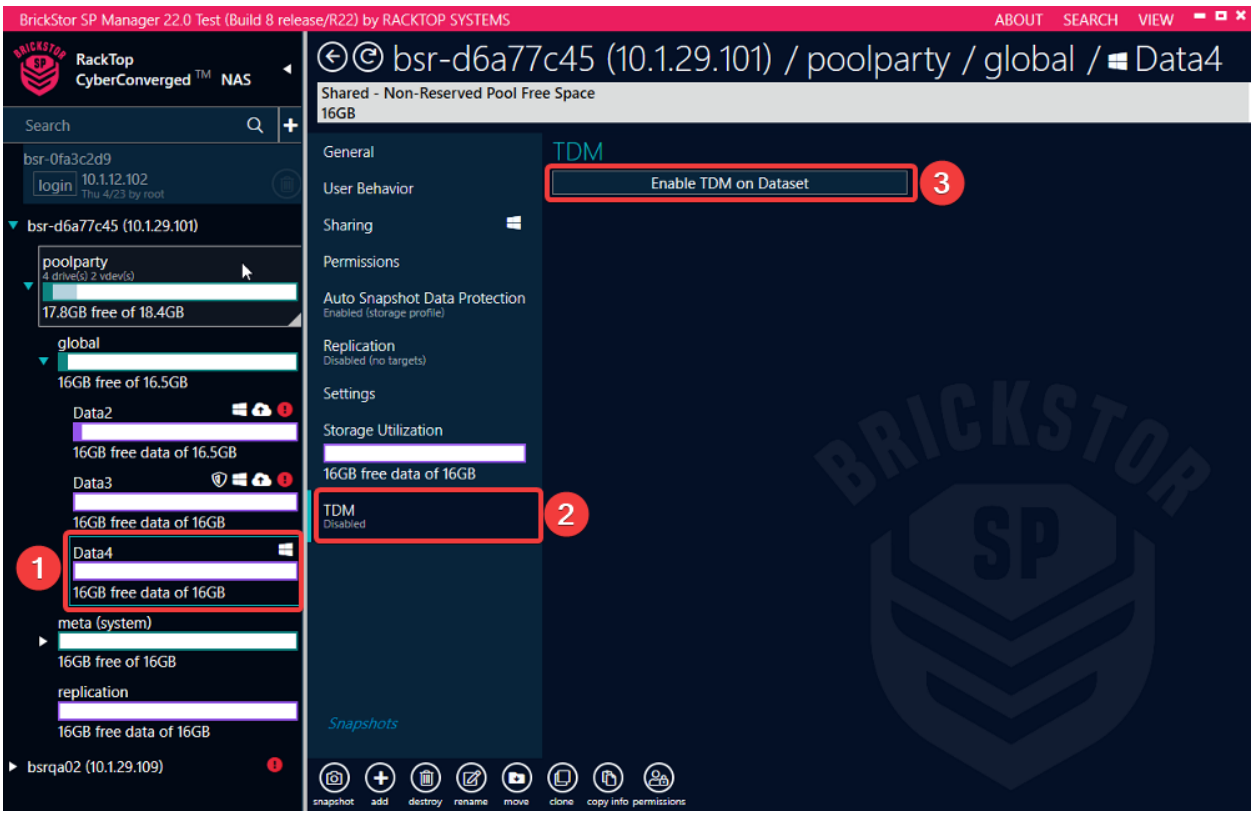

In the dialog, choose the location: New S3, NFS, or a location that has been set up previously, and choose the local space reclamation policy.

Local Space Reclamation

There are three ways that TDM will handle the local data once it has been uploaded to the remote location:

- None - Always keep local copy (mirror)
- Dynamic - Keep local copy until space is needed.
- Immediate - Remove local copy after upload.

In all cases, the data will appear to a client of a share to be local. When a request for a file is made, the file will be transparently downloaded from the remote location and returned to the client.

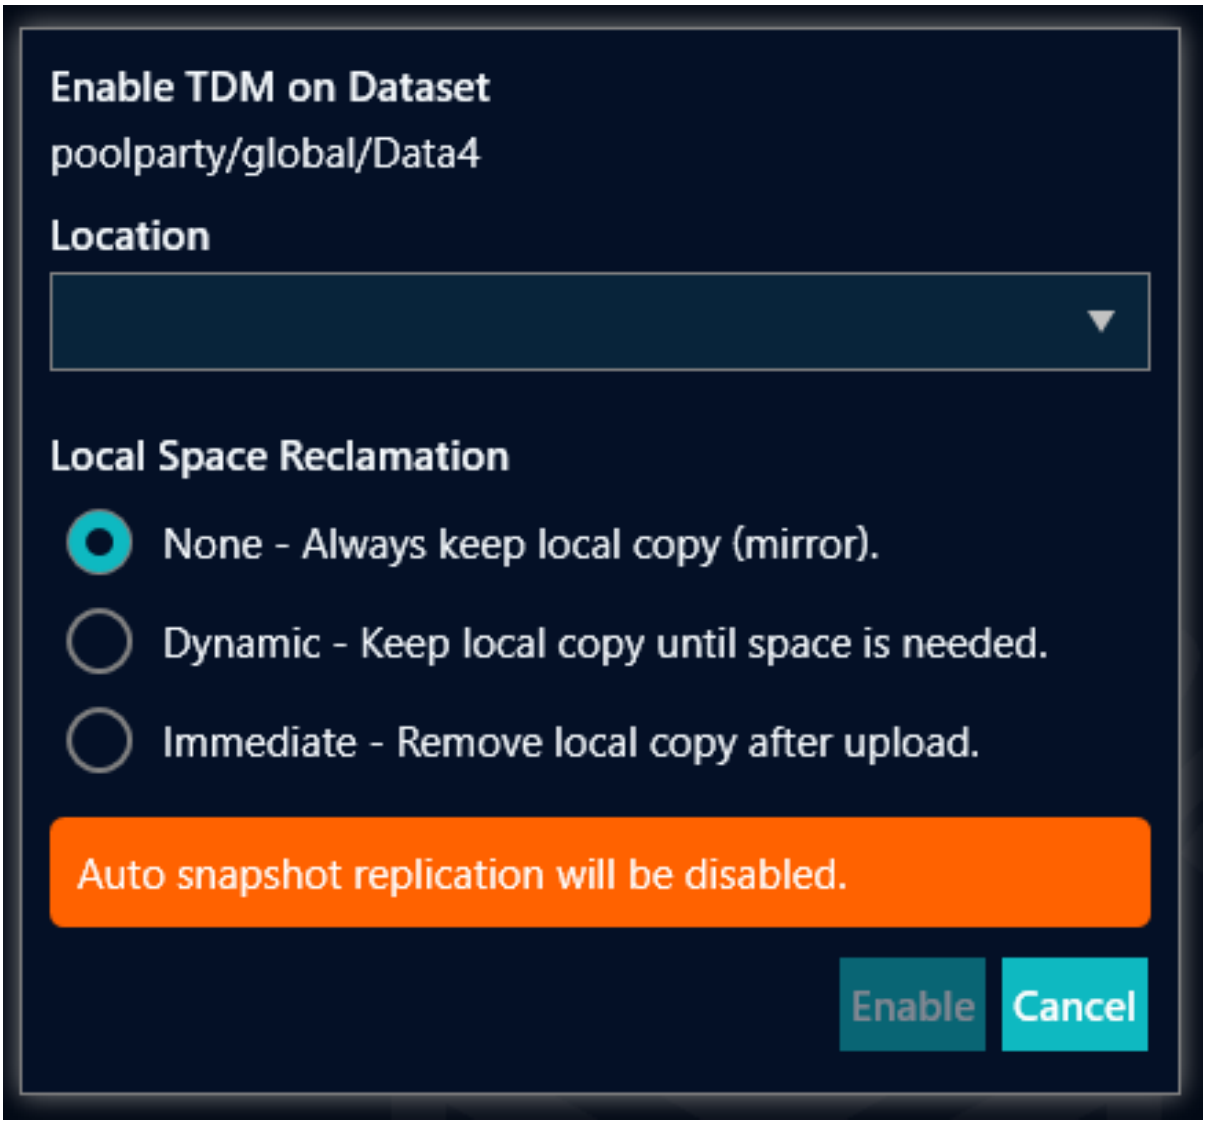

When configuring TDM to use an NFS location, provide the following details:

- Server
- Path

When configuring TDM to use an S3 bucket, provide the following details:

- Region
- Endpoint
- Bucket
- Object (optional)
- Access Key
- Secret Key

Example S3 configuration

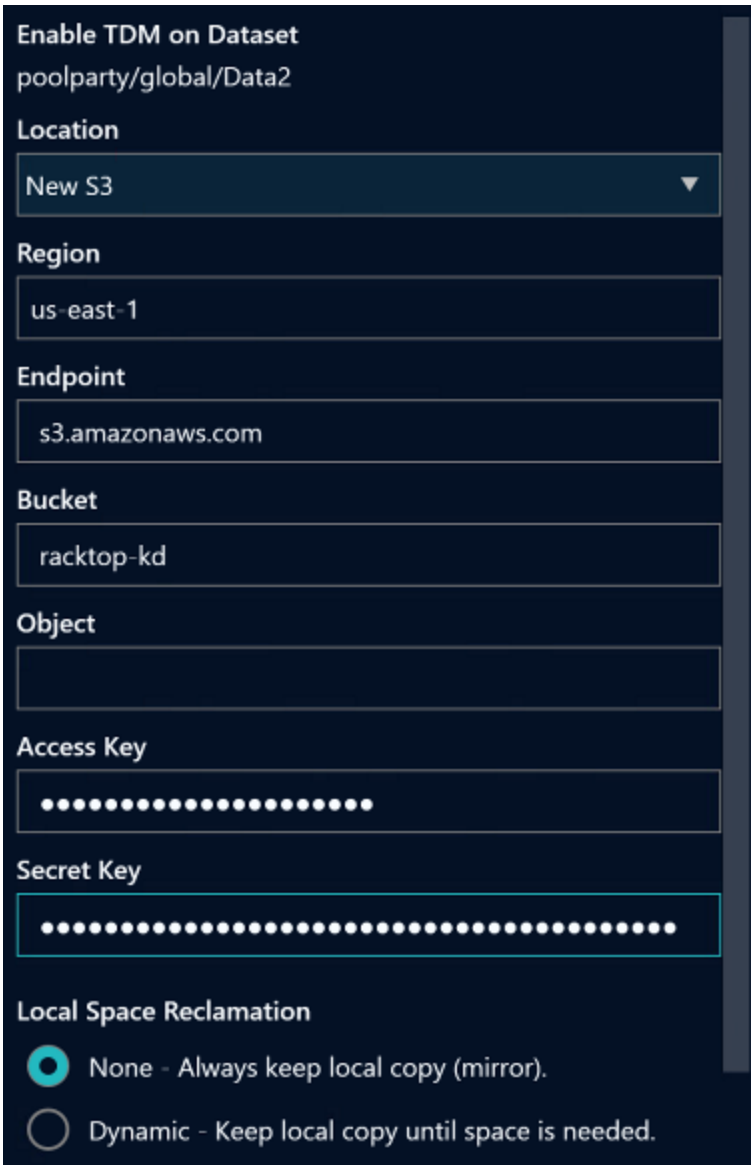

TDM Status and Data Distribution

In the dataset view, the Data Distribution graph will show the amount of data stored locally and the amount uploaded to the remote location.

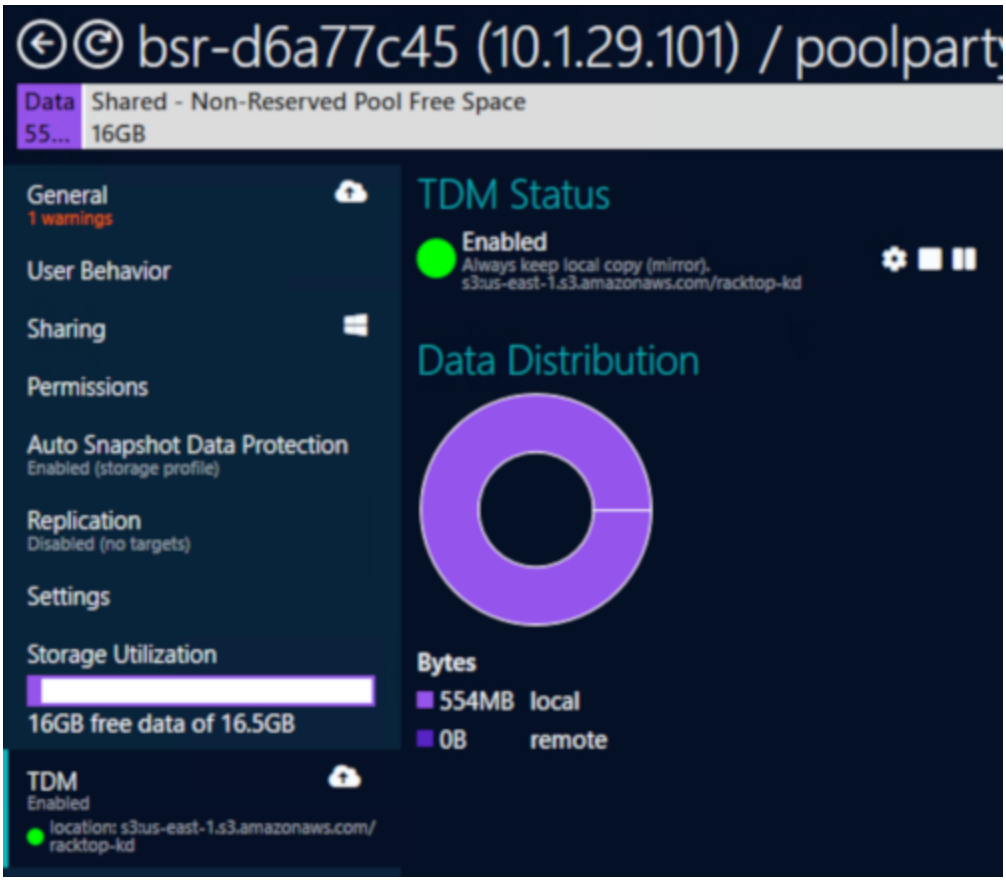

To view the status of TDM for all datasets or to reconfigure settings, choose the system in the Connections pane, and the TDM tab in the Detail pane.

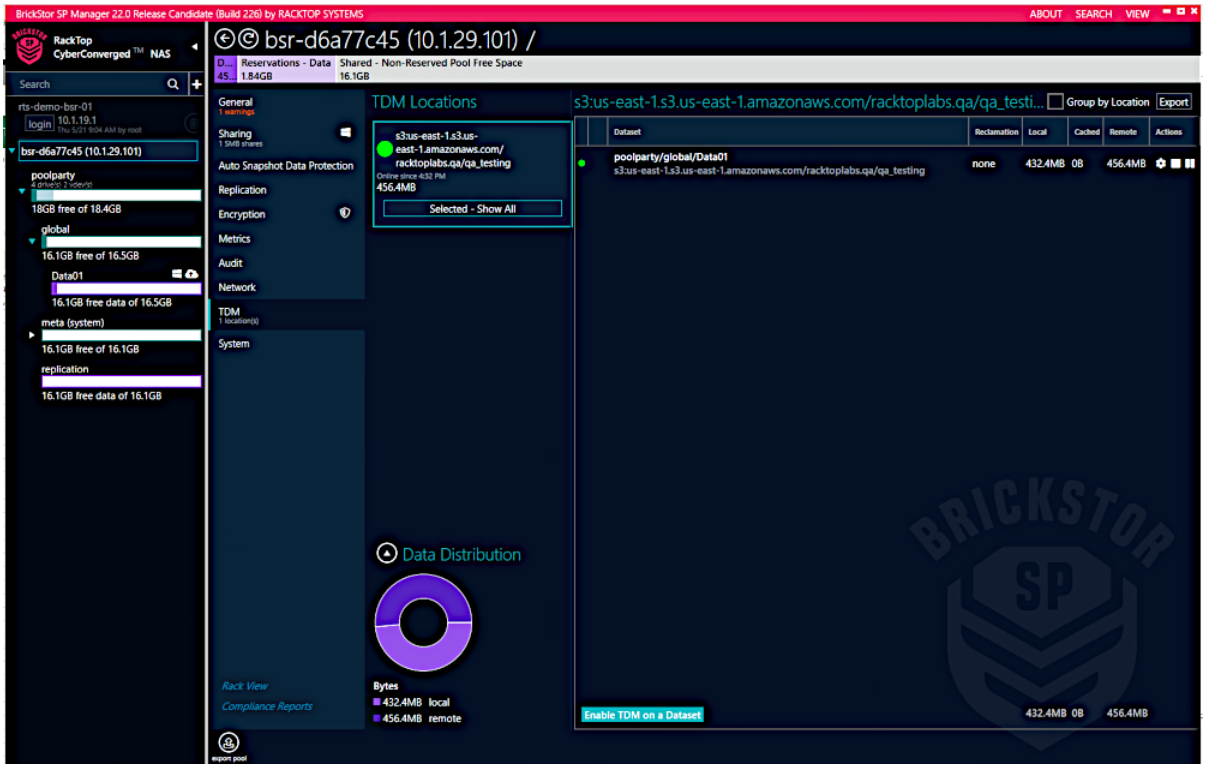

Reconfiguring TDM

Once TDM is set up, it is possible to reconfigure local space reclamation settings by clicking the gear icon:

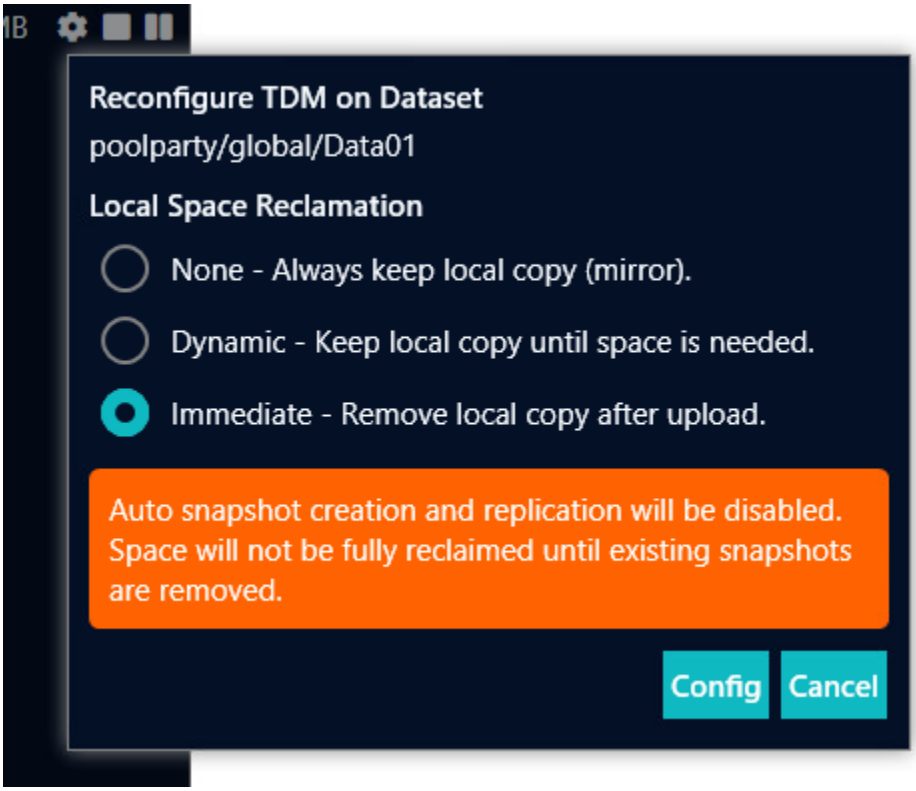

Disabling TDM

To disable TDM, click the stop icon.

The disable operation will download all remote files to the local dataset prior to disabling if local space reclamation was set to Immediate or Dynamic. You will be presented with a choice of whether to delete (remove), or leave (keep) the remote data.

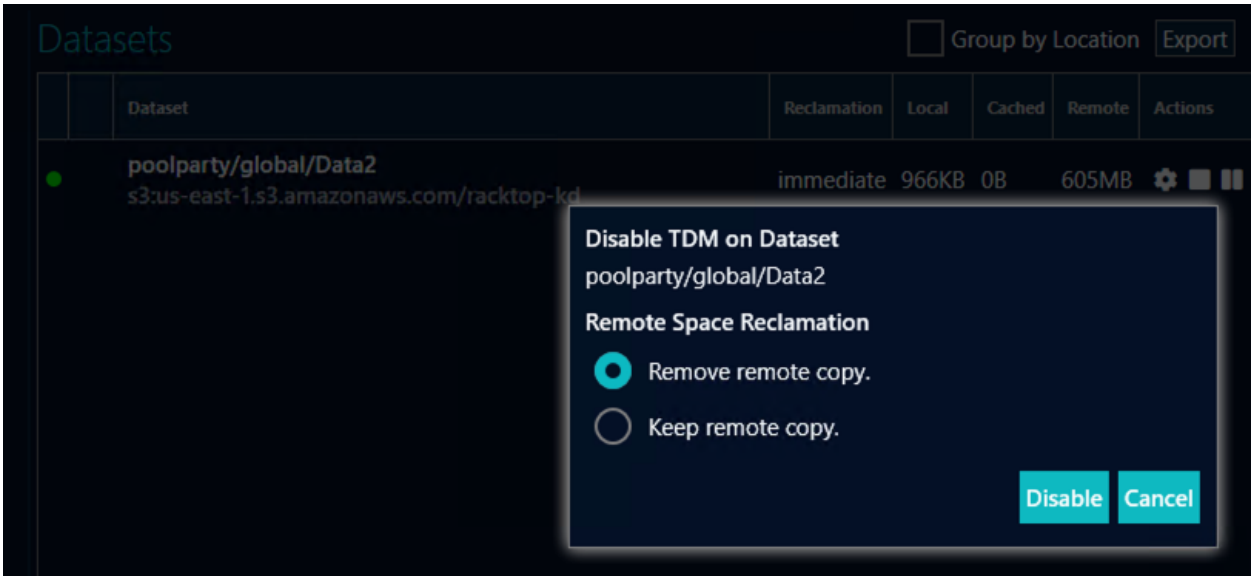

To complete the operation, acknowledge warnings and commit the change:

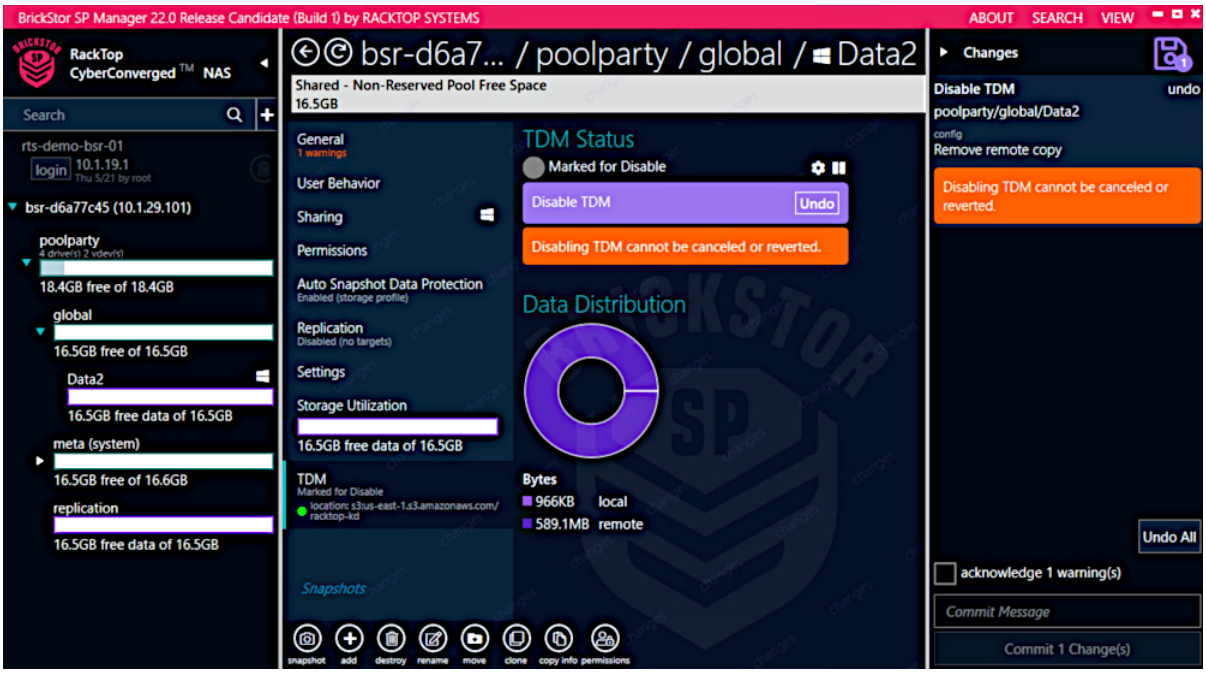

iSCSI Initiator

BrickStor supports connecting to iSCSI targets served from external NAS and other network-connected storage systems. Any connected iSCSI targets may be used in the same way as local storage to create new storage pools or to expand an existing one.

BrickStor can connect to iSCSI targets served from the following qualified third party systems:

- HPE Nimble AF40
- HPE 3Par 8400

Configuring the iSCSI Initiator

Complete the following steps to configure the BrickStor as an iSCSI Initiator:

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

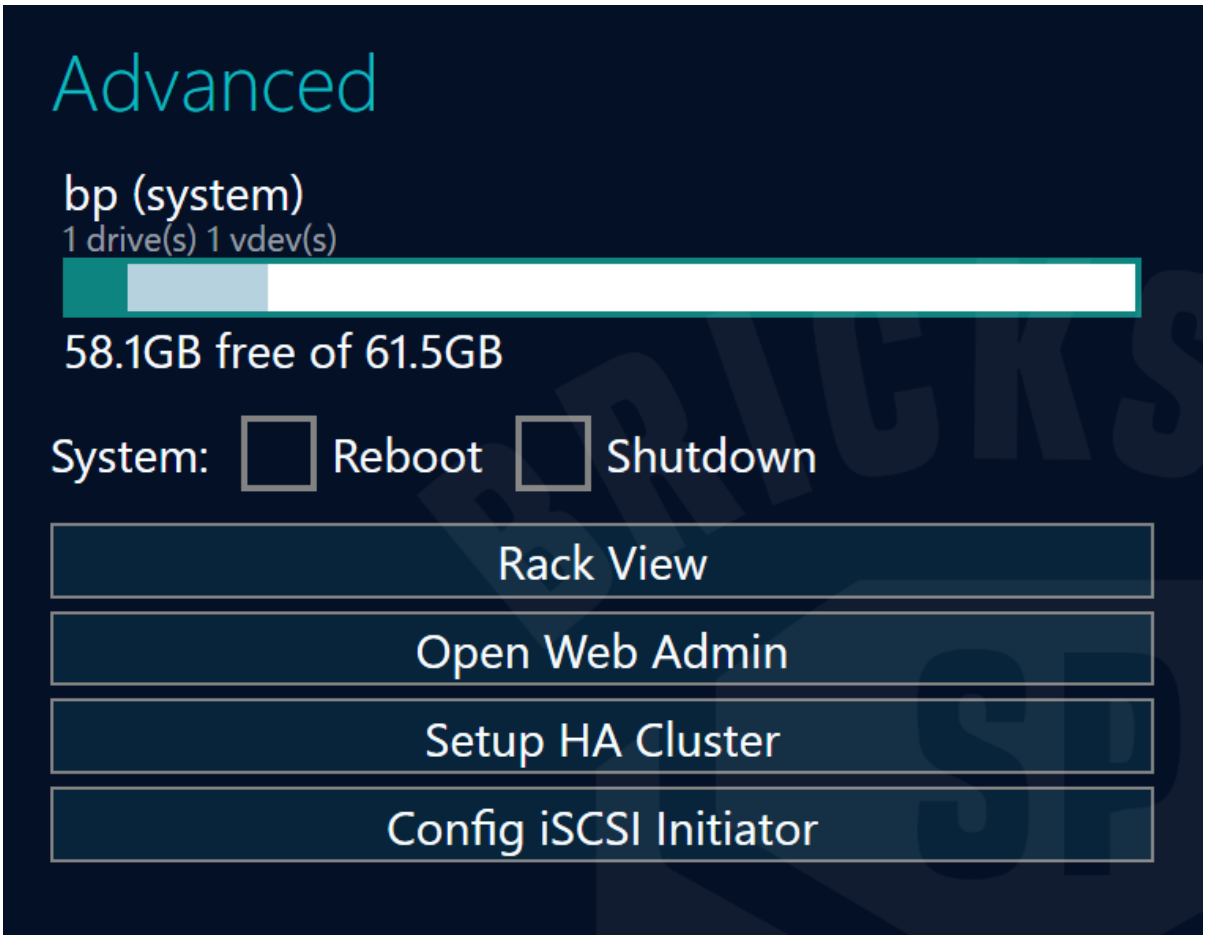

A new tab will be created in the Details Pane labeled **iSCSI Initiator**. Navigate to this tab to see the following information:

- The unique iSCSI Qualified Name (IQN) assigned to the Brickstor (i.e. `iqn.2010-03.io.racktop:zz00012k`)
- A button to **Configure Initiator**
- A button to **Add Target**

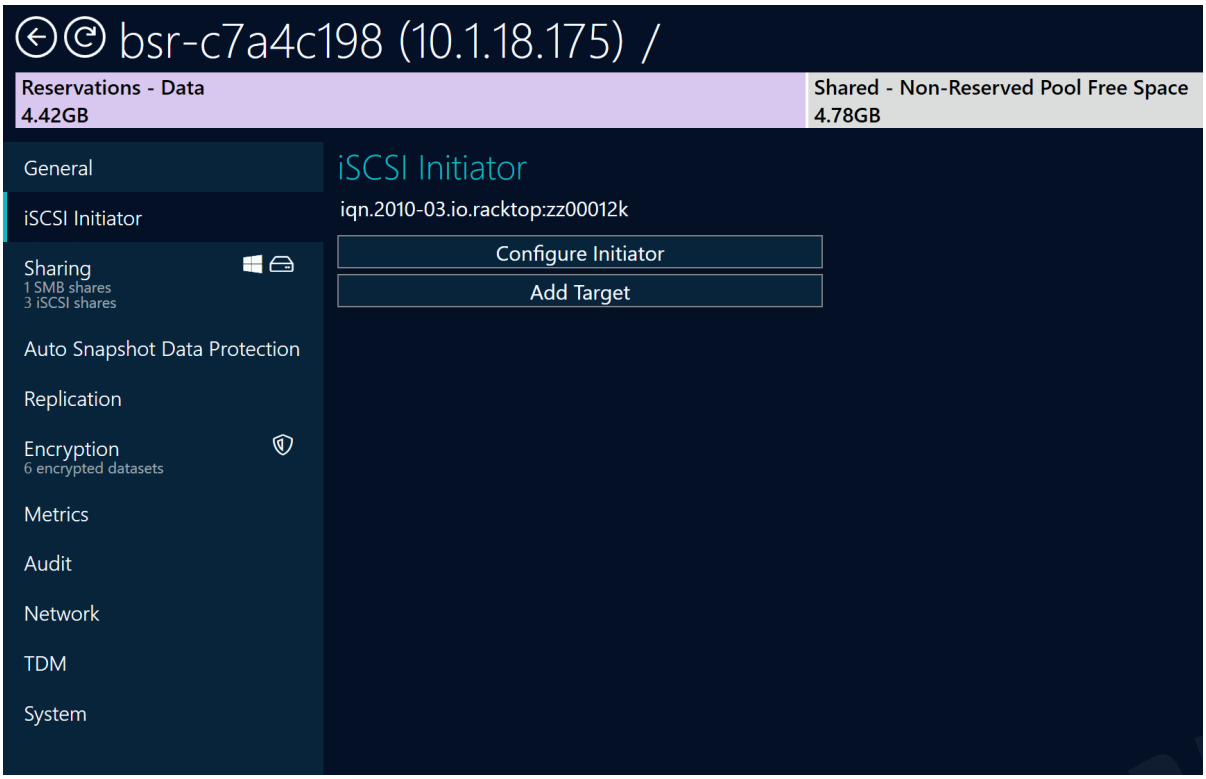

Configuring Initiator authentication

Once the **Configure Initiator** button is selected, a pop-up window appears where you can add a CHAP name and secret for 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 storage system that serves the target.

For **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 Initiator's CHAP Name, you can either use the already assigned IQN by selecting the associated **Use IQN** button or enter a free-text name.

For the Initiator's Chap Secret, you can enter a free text string. A minimum of 12 and a maximum of 16 characters are required.

Select **Apply** to apply the changes.

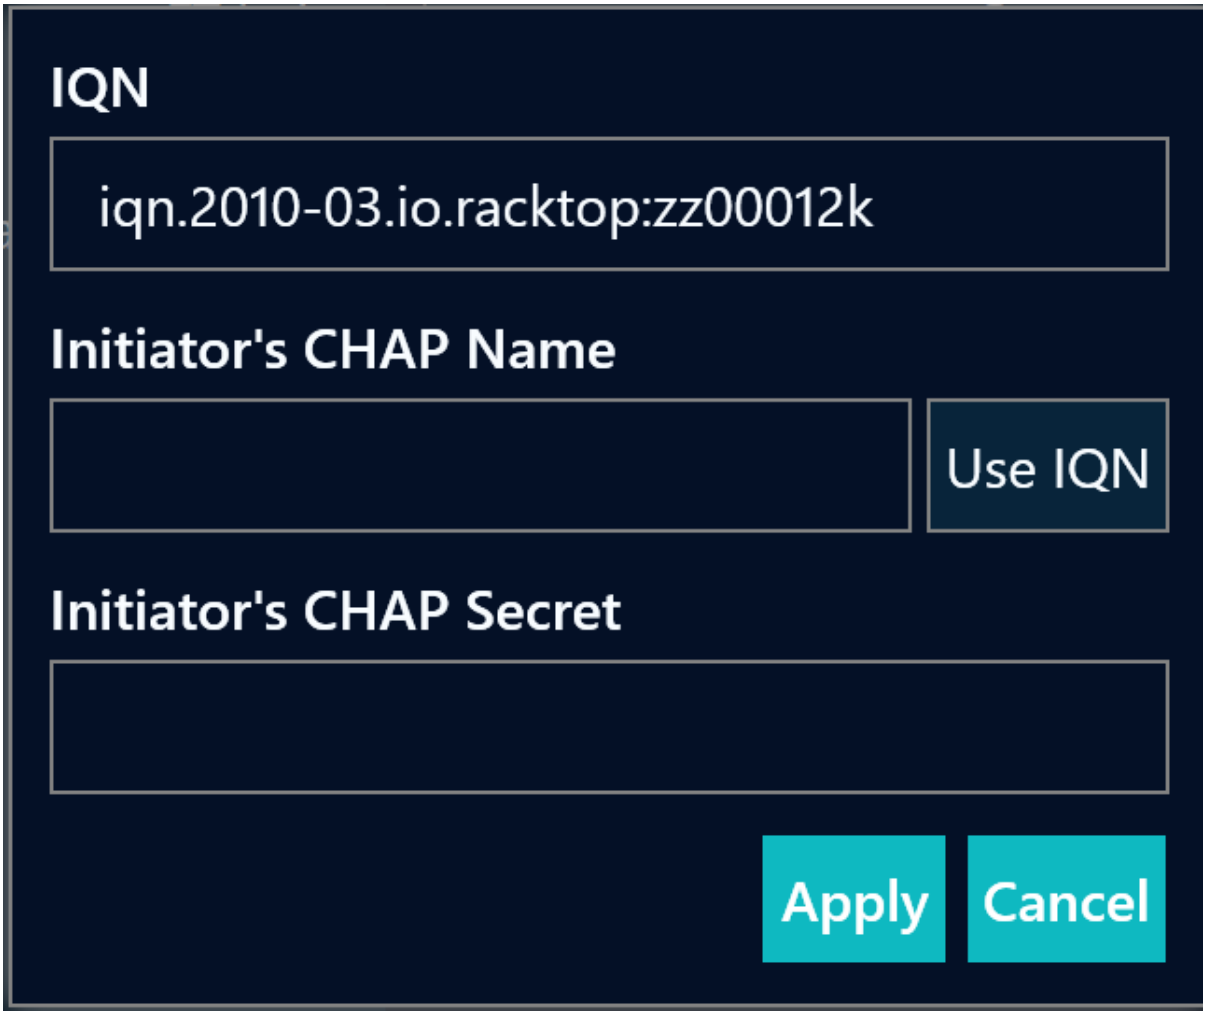

Connecting to the iSCSI Target

Once the **Add Target** button is selected, a pop-up window appears where you can add an iSCSI Target to BrickStor. Adding an iSCSI Target will make it available as a block, or disk, device. Such devices can be used to create new storage pools on BrickStore or to expand an existing pool.

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 can 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 **TPG Tag**, a numeric value may be specified which corresponds to a Target Portal Group (TPG) on the Target.

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

If enabling a two-way CHAP is desired, you can elect to **Enable Chap**. Doing so will require you to enter the following for the Target:

- Mutual CHAP - Target's Name
- Mutual CHAP - Target's Secret

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

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, volumes that are associated with it will be shown as being offline.

NOTE The configuration setting to the Target can be modified, as well as the ability to remove the iSCSI Target altogether, using the associated icons.

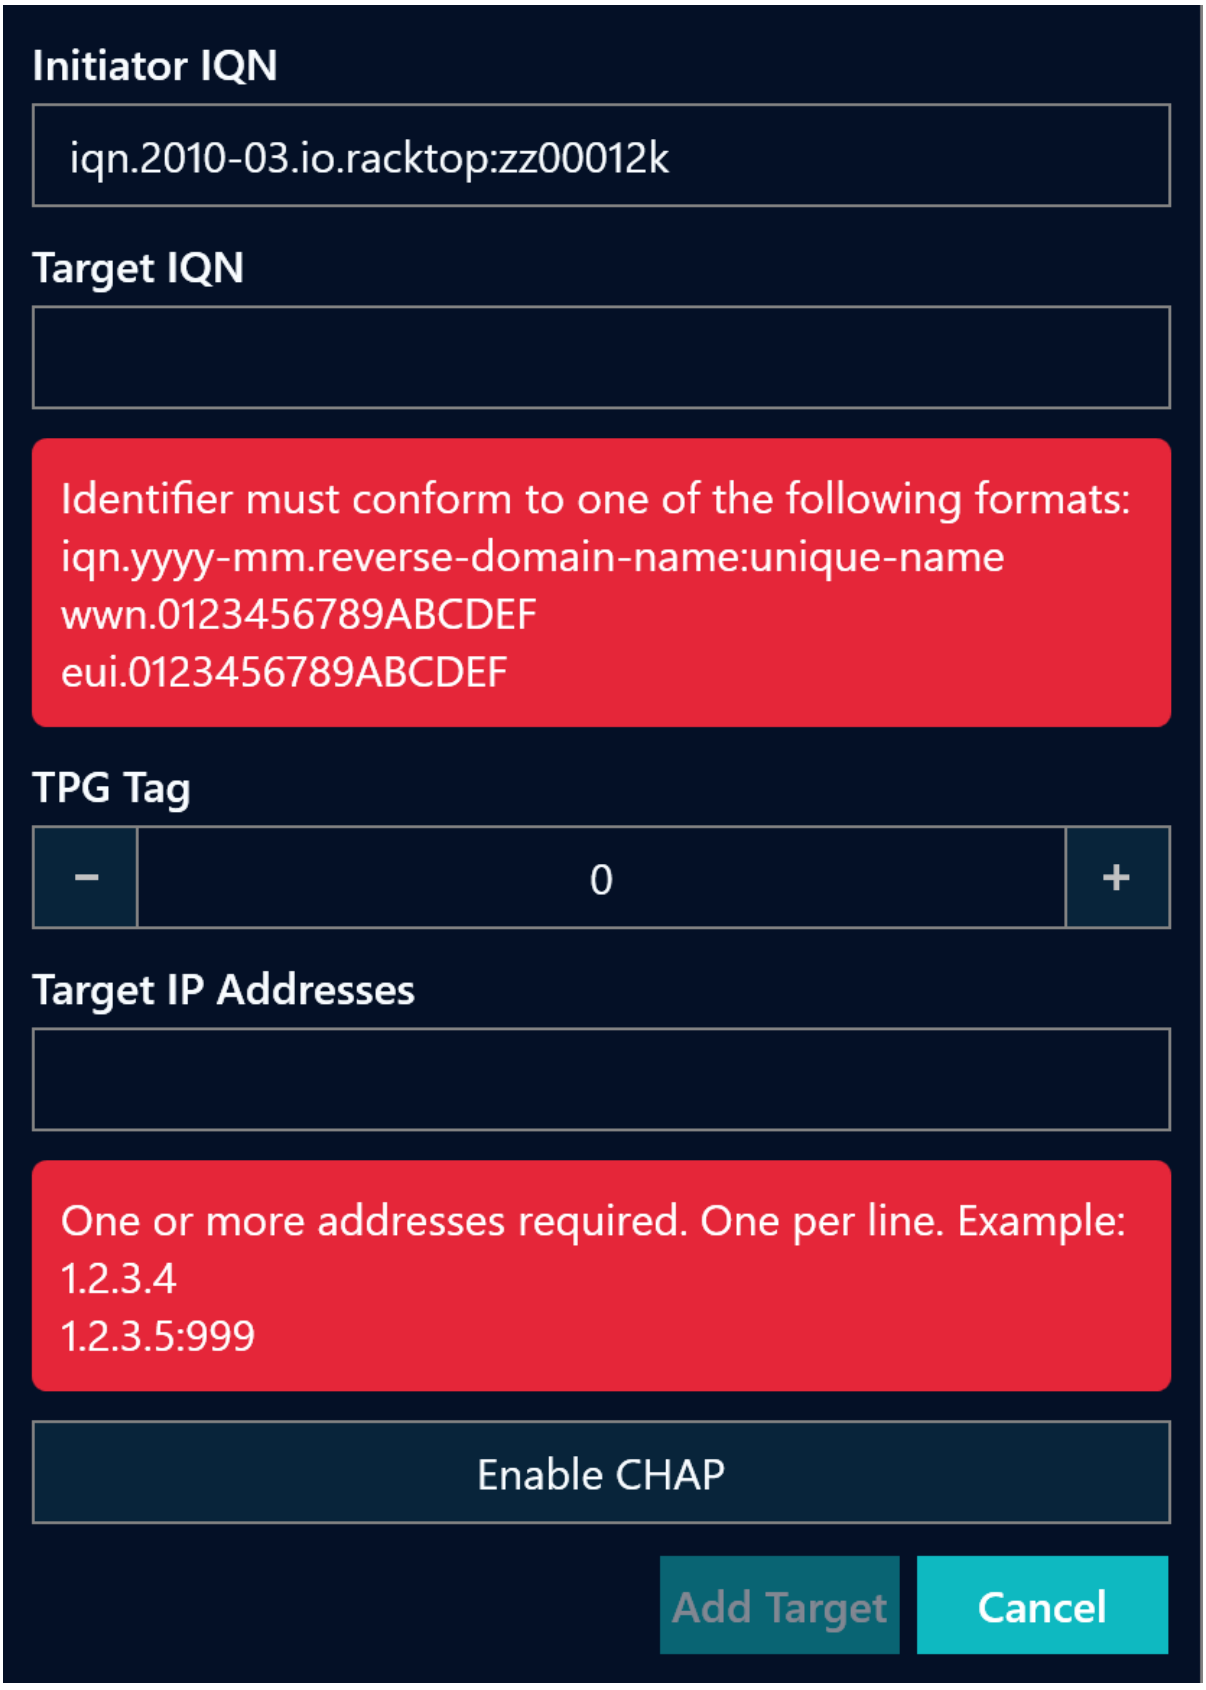
Initiator IQN
iqn.2010-03.io.racktop:zz00012k
Target IQN
Identifier must conform to one of the following formats:
iqn.yyyy-mm.reverse-domain-name:unique-name
wwn.0123456789ABCDEF
eui.0123456789ABCDEF
TPG Tag
−
0
+
Target IP Addresses
One or more addresses required. One per line. Example:
1.2.3.4
1.2.3.5:999
Enable CHAP
Add Target
Cancel

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.

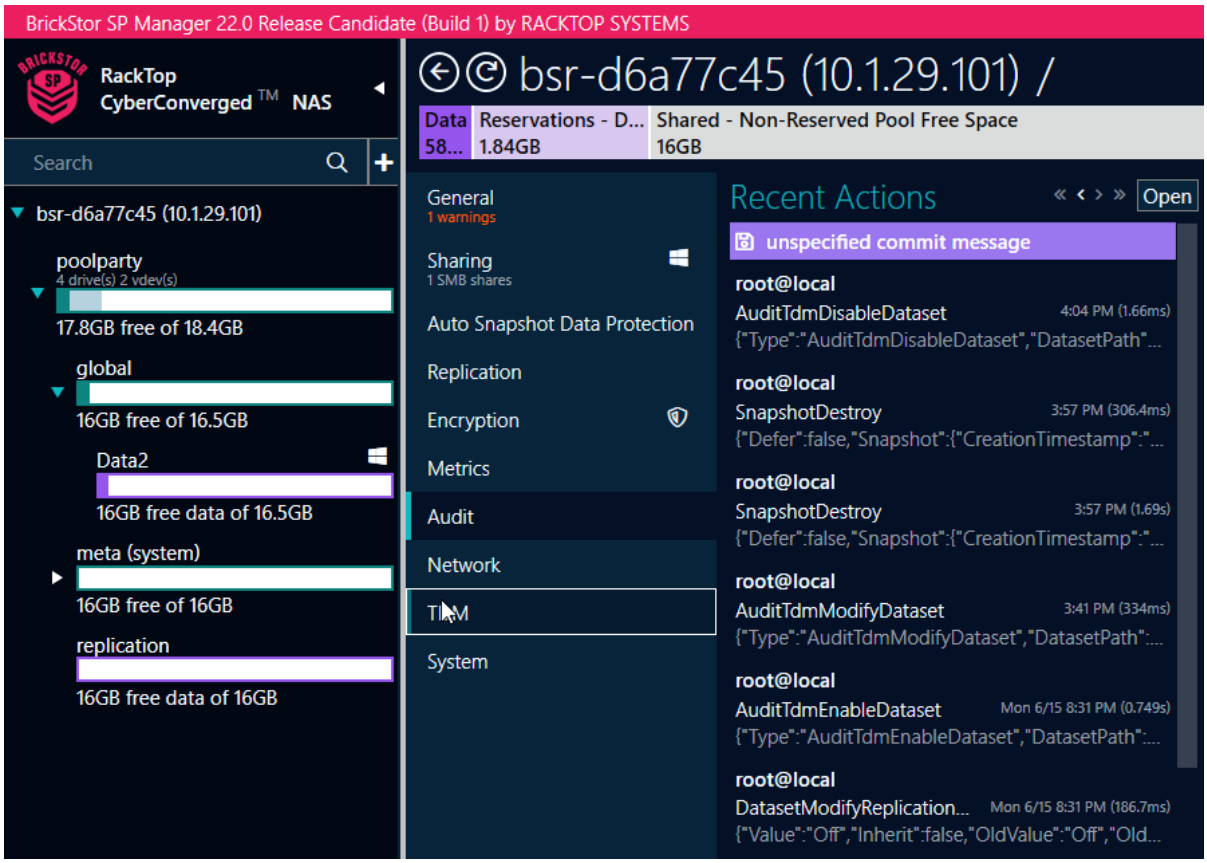

3. Hover your pointer over any of the actions to display all of the API messages posted for the change.

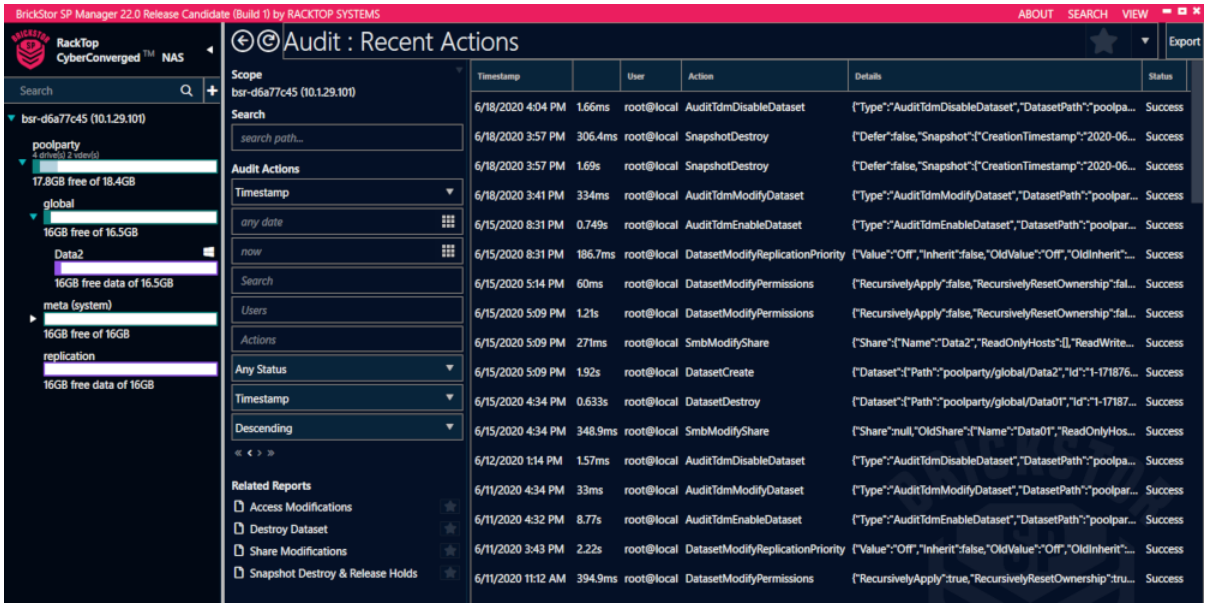
BrickStor SP Manager 22.0 Release Candidate (Build 1) by RACKTOP SYSTEMS
ABOUT SEARCH VIEW
RackTop CyberConverged NAS
Search
bsr-d6a77c45 (10.1.29.101)
poolparty
17.8GB free of 18.4GB
global
16GB free of 16.5GB
Data2
16GB free data of 16.5GB
meta (system)
16GB free of 16GB
replication
16GB free data of 16GB
Audit : Recent Actions
Export
Scope
bsr-d6a77c45 (10.1.29.101)
Search
search path...
Audit Actions
Timestamp
any date
now
Search
Users
Actions
Any Status
Timestamp
Descending
Related Reports
Access Modifications
Destroy Dataset
Share Modifications
Snapshot Destroy & Release Holds
Timestamp User Action Details Status
6/18/2020 4:04 PM 1.66ms root@local AuditTdmDisableDataset {"Type":"AuditTdmDisableDataset","DatasetPath":"poolpa... Success
6/18/2020 3:57 PM 306.4ms root@local SnapshotDestroy {"Defer":false,"Snapshot":{"CreationTimestamp":"2020-06... Success
6/18/2020 3:57 PM 1.69s root@local SnapshotDestroy {"Defer":false,"Snapshot":{"CreationTimestamp":"2020-06... Success
6/18/2020 3:41 PM 334ms root@local AuditTdmModifyDataset {"Type":"AuditTdmModifyDataset","DatasetPath":"poolpar... Success
6/15/2020 8:31 PM 0.749s root@local AuditTdmEnableDataset {"Type":"AuditTdmEnableDataset","DatasetPath":"poolpar... Success
6/15/2020 8:31 PM 186.7ms root@local DatasetModifyReplicationPriority {"Value":"Off","Inherit":false,"OldValue":"Off","OldInherit":... Success
6/15/2020 5:14 PM 60ms root@local DatasetModifyPermissions {"RecursivelyApply":false,"RecursivelyResetOwnership":fal... Success
6/15/2020 5:09 PM 1.21s root@local DatasetModifyPermissions {"RecursivelyApply":false,"RecursivelyResetOwnership":fal... Success
6/15/2020 5:09 PM 271ms root@local SmbModifyShare {"Share":{"Name":"Data2","ReadOnlyHosts":[],"ReadWrite... Success
6/15/2020 5:09 PM 1.92s root@local DatasetCreate {"Dataset":{"Path":"poolparty/global/Data2","Id":"1-171876... Success
6/15/2020 4:34 PM 0.633s root@local DatasetDestroy {"Dataset":{"Path":"poolparty/global/Data01","Id":"1-17187... Success
6/15/2020 4:34 PM 348.9ms root@local SmbModifyShare {"Share":null,"OldShare":{"Name":"Data01","ReadOnlyHos... Success
6/12/2020 1:14 PM 1.57ms root@local AuditTdmDisableDataset {"Type":"AuditTdmDisableDataset","DatasetPath":"poolpa... Success
6/11/2020 4:34 PM 33ms root@local AuditTdmModifyDataset {"Type":"AuditTdmModifyDataset","DatasetPath":"poolpar... Success
6/11/2020 4:32 PM 8.77s root@local AuditTdmEnableDataset {"Type":"AuditTdmEnableDataset","DatasetPath":"poolpar... Success
6/11/2020 3:43 PM 2.22s root@local DatasetModifyReplicationPriority {"Value":"Off","Inherit":false,"OldValue":"Off","OldInherit":... Success
6/11/2020 11:12 AM 394.9ms 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.

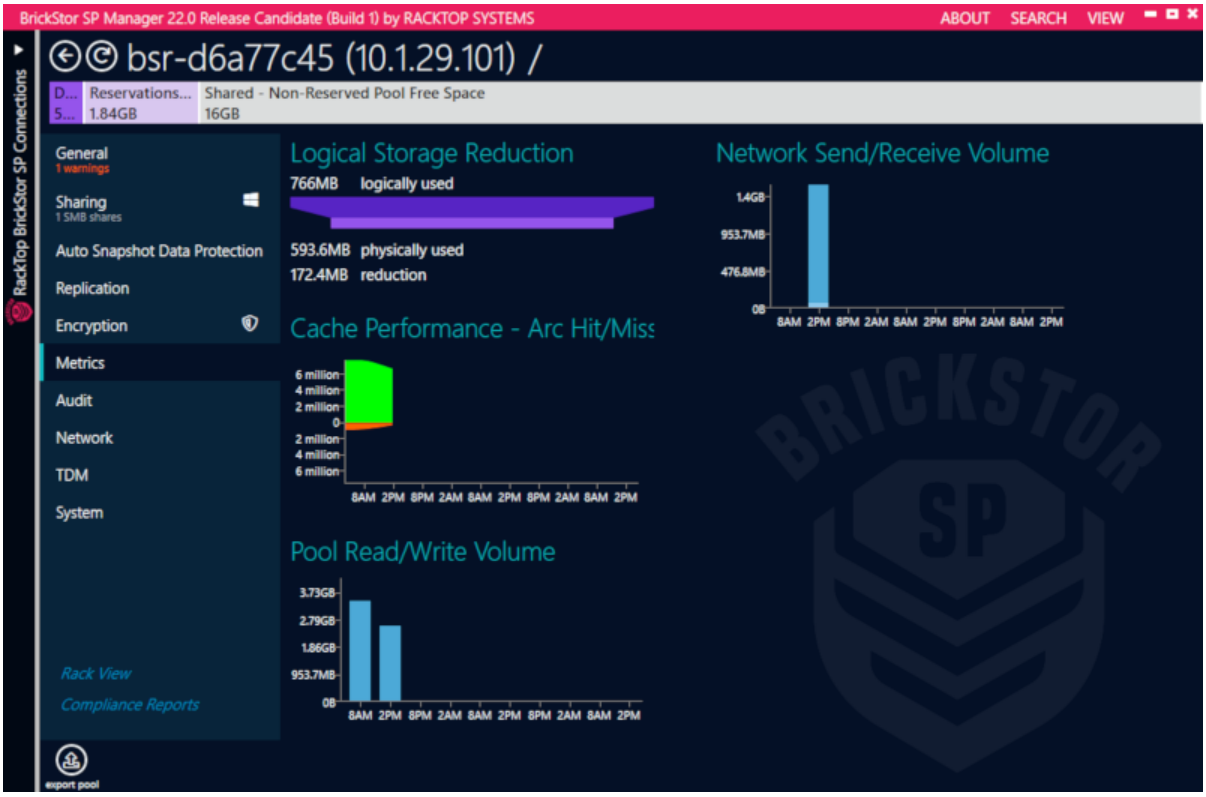

Accessing Metrics

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.

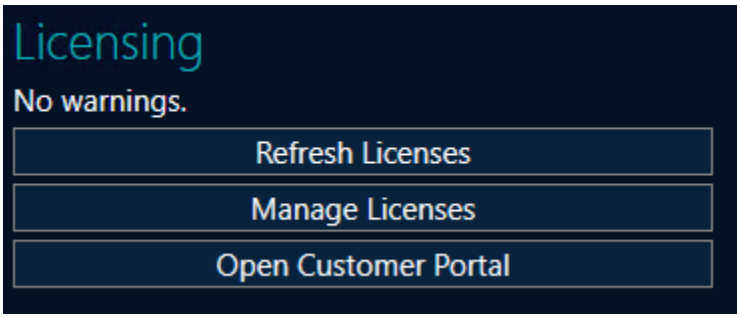

Refresh Licenses

Systems connected to the internet will automatically retrieve newly assigned or updated licenses from the MyRackTop customer portal. To pull down and apply licenses now click the **Refresh Licenses** button.

Manage Licenses

Manage Licenses will open a web browser to the managed BrickStor SP HTML5 user interface. Once logged in, it will display currently applied licenses and allow adding new ones.

To add a new license, enter a license key and click the **Add Key** button.

NOTE Applying subscription licenses such as Maintenance will replace an existing, expiring license.

Open Customer Portal

The Open Customer Portal button will open a web browser to the MyRackTop Customer portal to view details about this BrickStor SP system.

Health

The Heath Tab shows the health of the system. Probes represent the health of each component on the system. A component can have one or more probes representing various health aspects of the component. For example, each pool will have a probe for the status of the pool as well as a probe for the capacity of the pool. When a probe detects an issue, it will create alarms that can be viewed in the Health Tab. Each alarm has an associated severity. Currently, alarm severities include **Warning**, **Error**, and **Critical**.

The behavior of an alarm depends on the type of probe that generated the alarm. There are currently two types of probes. The first (and most common) type of probe is a sensor. Sensors generate alarms based on a measured value. Alarms caused by sensor probes will automatically clear themselves once the condition that caused the alarm has ceased.

The other type of probe is a Log probe. Log probes generate alarms based on values observed from a log file on the system. Unlike sensor probes, alarms generated by log probes do not resolve themselves. Instead, they must be explicitly acknowledged by the operator.

Accessing the Health Tab

To access the Health tab, complete the following steps:

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

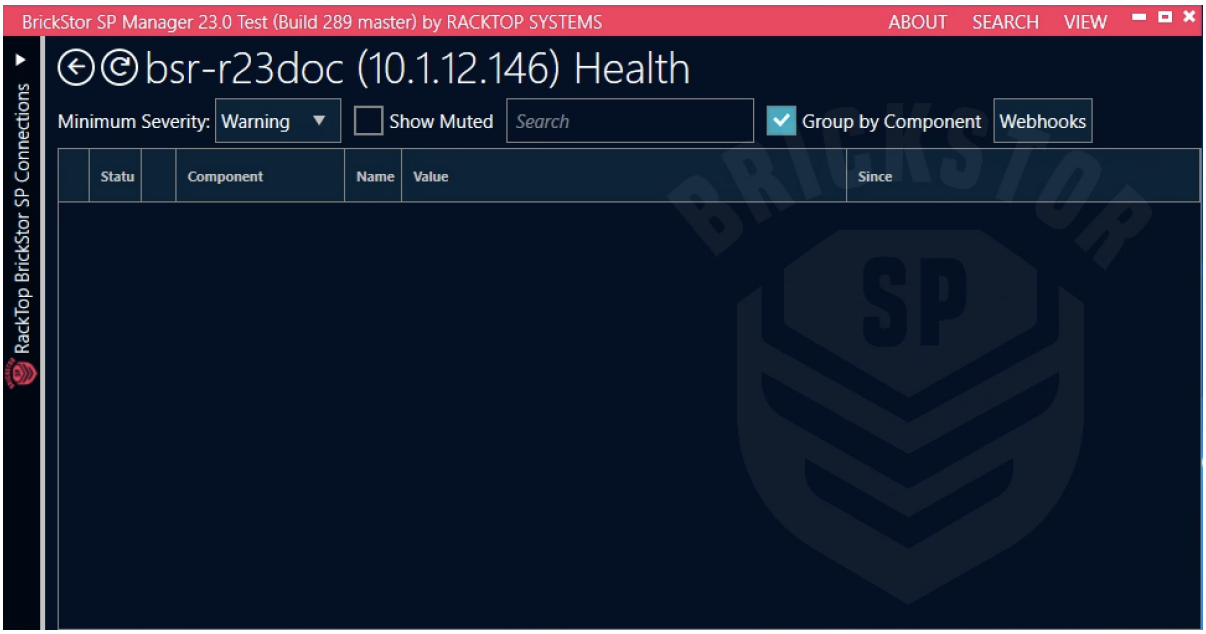

By default, probes are grouped by the component associated with the probe. To show each probe individually, uncheck the **Group by Component** checkbox at the top of the Health tab.

The list of probes can also be filtered by alarm severity or name. Clicking on the **Minimum Severity** dropdown allows you to change the minimum alarm severity (or select **All** to show all probes, including those without any alarms). To filter by a name, enter the name of the probe in the **Search** box. To show muted probes, check the **Show Muted** checkbox.

Health Tasks

Showing Probe Details

To show probe details, locate the probe in the probe table. Click on the ▶ symbol to expand the probe details. Clicking on the ▼ symbol will collapse the probe details.

Clicking on the **Details** button will open a window displaying a JSON definition of the probe (including the event history). To copy the probe details as JSON to the clipboard, click the **Copy** button.

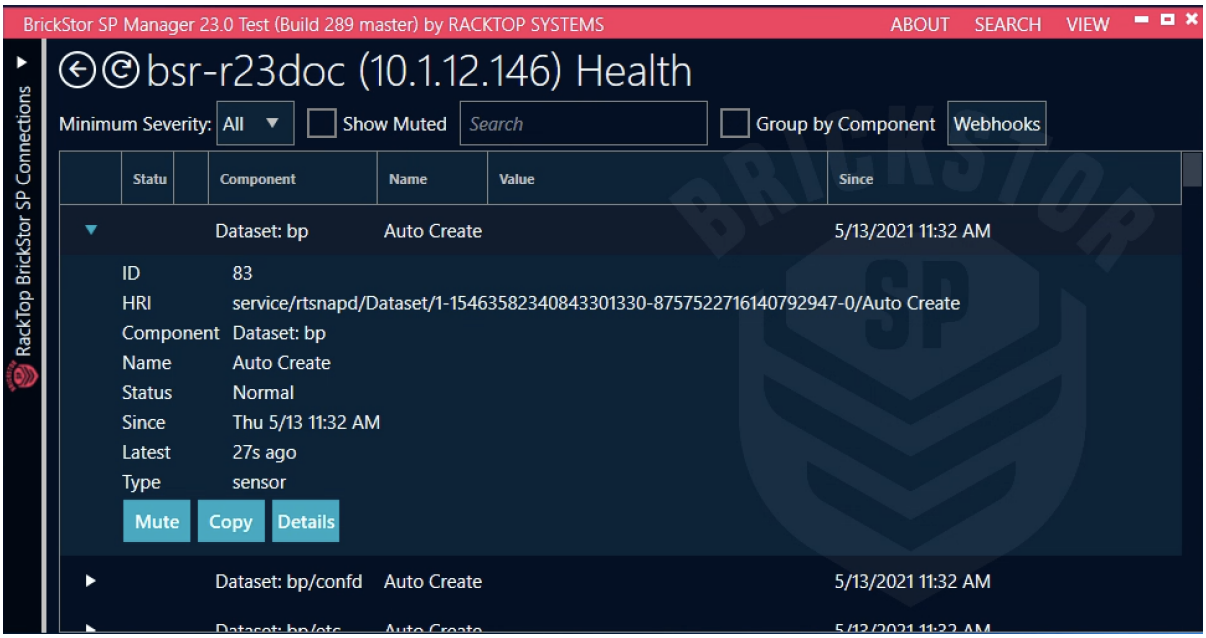

Acknowledging an Alarm

Unlike sensor-based probes, log-based probes must be acknowledged to clear an alarm. To acknowledge an alarm, complete the following steps:

1. Locate the probe corresponding to the alarm in the probe table.
2. Click on the gear (⚙) icon besides the component name.
3. Select **Ack**.

Muting a Probe

Muting a probe prevents webhooks from being invoked as well as prevents email alerts from being sent. To mute a probe, complete the following steps:

1. Locate the probe to mute in the probe table.
2. Click on the gear icon (⚙) beside the component name.
3. Select Mute
4. Select the desired duration to mute the probe. This can be forever (until explicitly unmuted) or for

a duration of up to a year.

5. By default, the probe will be unmuted if the severity of the probe changes. To disable this behavior, uncheck the **Unmute if severity changes** box.
6. Click **Mute**

Alternatively, one can expand the probe details and click the **Mute** button from the expanded probe details.

Pruning A Stale Probe

When a probe stops reporting data, it becomes stale. This is a rare occurence and normally shouldn't be encountered. When this does occur, a **Prune** option will appear when clicking the gear (⚙) icon. Clicking on the **Prune** option will remove the stale probe.

Webhooks

BrickStor SP can be configured to emit Webhook notifications. A Webhook is a way to send a varety of events to an external service which implements a Webhook API. This event may or may not be due to direct action from BrickStor SP or the Manager.

The Webhook configuration page can be accessed via the **System** tab. On that tab, in the **Advanced** section, select **Webhooks** to begin configuration.

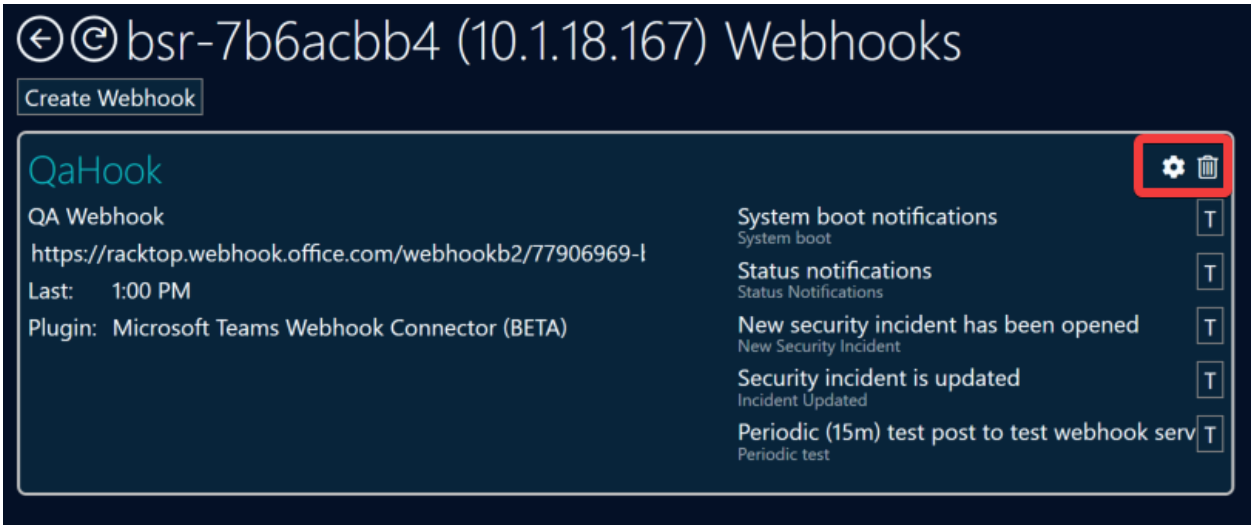

The are several types of Webhooks available for creation. The webhooks include:

- Microsoft Teams Webhook Connector
- Pager Duty Notifications
- Pushover Notifications
- Slack Webhook Connector
- Racktop Webhook Format

The next sections include examples of how to configure each of these Webhooks services.

Microsoft Teams Webhook Connector

This Webhook allows you to send BrickStor SP notifications to the Microsoft Teams Application. First ensure you have performed the necessary steps to establish Webhook connections within Microsoft Teams.

Next, do the following on BrickStor SP:

1. Select **Create Webhook** on the **System** tab.
2. Enter a **Name** for the Webhook
3. Enter a **Description** for the Webhook
4. For **Type** select **Microsoft Teams Webhook Connector**
5. For **URL**, paste in the URL provided from the Teams application.

6. Select to turn on / off debugging to include event metadata.
7. Choose the event types to which you would like to subscribe. You may select any number of event types
8. Select **Create**

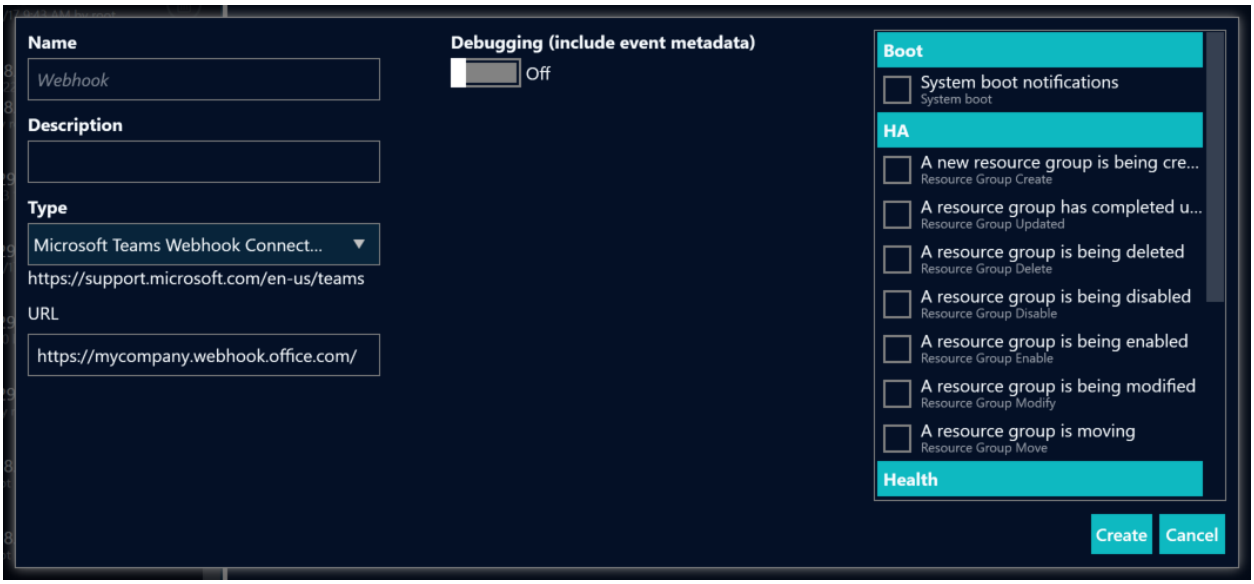

Pager Duty Notifications

This Webhook allows you to send BrickStor SP notifications to the PagerDuty incident response platform. First ensure you have performed the necessary steps to establish Webhook connections within PagerDuty.

Next, do the following on BrickStor SP:

1. Select **Create Webhook** on the **System** tab.
2. Enter a **Name** for the Webhook
3. Enter a **Description** for the Webhook
4. For **Type**, select **PagerDuty Notifications**
5. For **URL**, paste in the URL provider from the PagerDuty platform.
6. Enter the required Integration Key, also known as the Routing key.
7. Optionally, you can enter a changed URL derived from the main URL as an advanced option.
8. Choose the event types to which you would like to subscribe. You may select any number of event types
9. Select **Create**

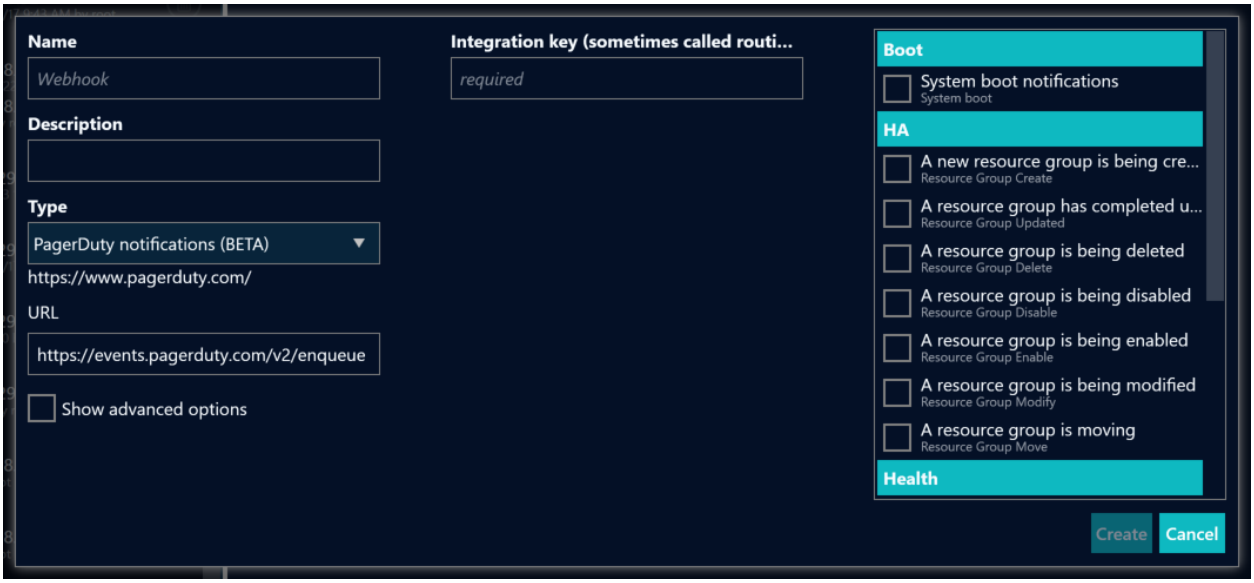

Pushover Notifications

This Webhook allows you to send BrickStor SP notifications to the Pushover App for real-time notifications on your smart device. First ensure you have performed the necessary steps to establish Webhook connections within Pushover.

Next, do the following on BrickStor SP:

1. Select **Create Webhook** on the **System** tab.
2. Enter a **Name** for the Webhook
3. Enter a **Description** for the Webhook
4. For **Type**, select **Pushover Notifications**
5. For **URL**, paste in the URL provider from the Pushover App.
6. Enter the required **Pushover API token**.
7. Enter the **Users/Groups** to receive notification.
8. Enter the **Devices** to send notifications.
9. Optionally, you can select the override default priority and sounds to play on incoming alerts and emergencies.
10. Choose the event types to which you would like to subscribe. You may select any number of event types
11. Select **Create**

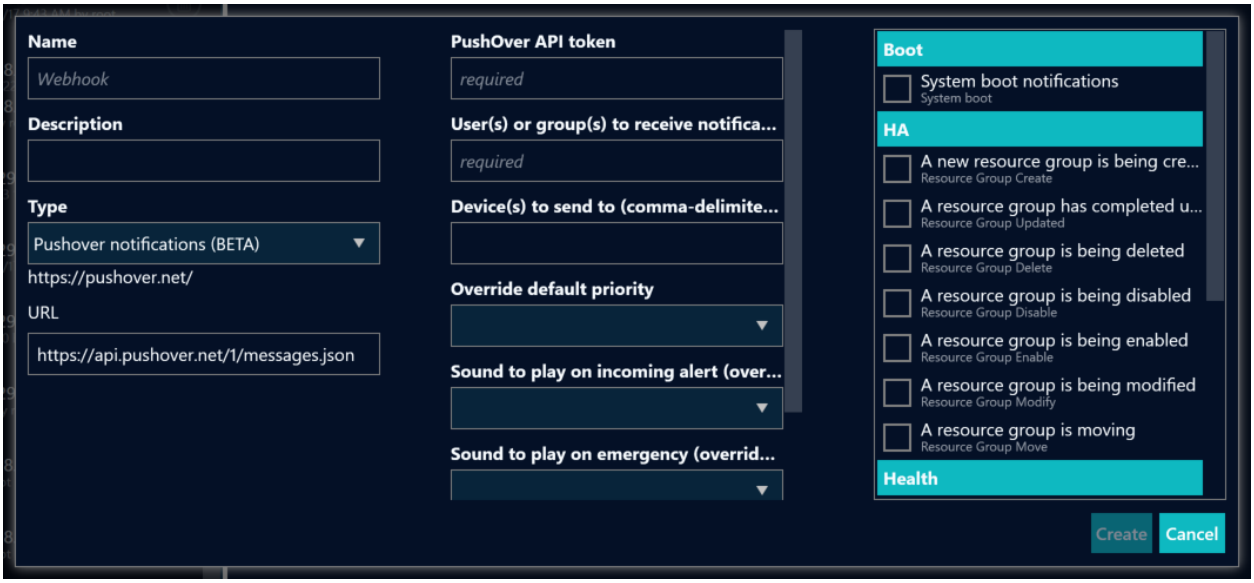

Slack Webhook Connector

This Webhook allows you to send BrickStor SP notifications to the Slack App for real-time notifications on your smart device. First ensure you have performed the necessary steps to establish Webhook connections within Slack.

Next, do the following on BrickStor SP:

1. Select **Create Webhook** on the **System** tab.
2. Enter a **Name** for the Webhook
3. Enter a **Description** for the Webhook
4. For **Type**, select **Slack Webhook Connector**
5. For **URL**, paste in the URL provider from the Slack App.
6. Choose the event types to which you would like to subscribe. You may select any number of event types
7. Select **Create**

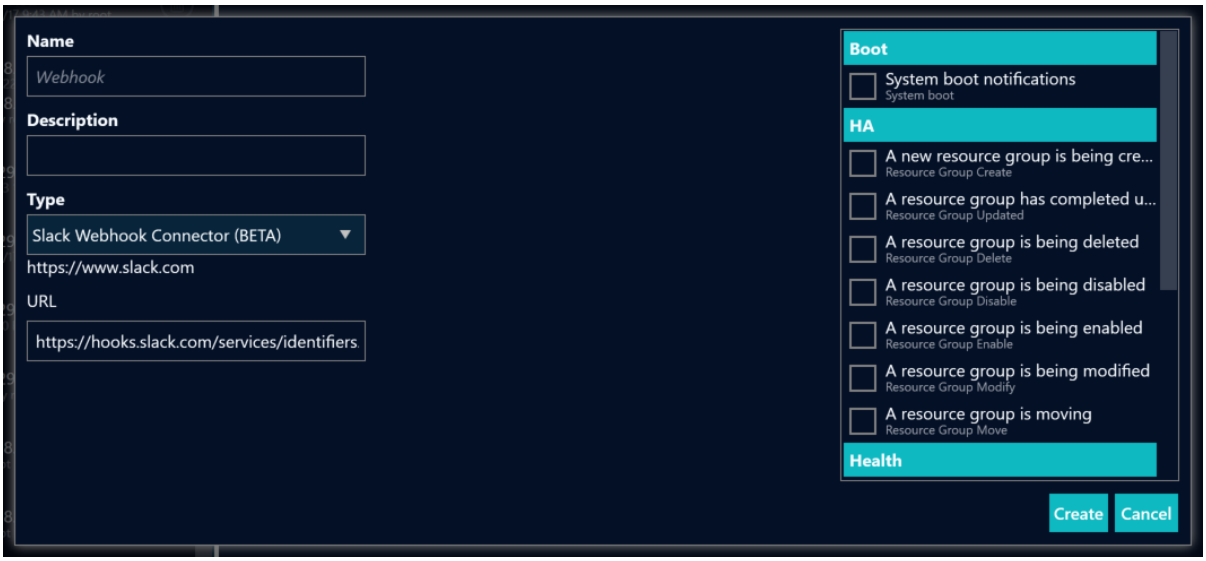

RackTop Webhook Format

This Webhook allows the user to use a generic connection to attempt to send BrickStor SP notifications to other applications not listed

Next, do the following on BrickStor SP:

1. Select **Create Webhook** on the **System** tab.
2. Enter a **Name** for the Webhook
3. Enter a **Description** for the Webhook
4. For **Type**, select **Racktop Webhook Format**
5. For **URL**, paste in the URL provider from the external application.
6. Optionally, you can provide a Username / Password / Secret for basic HTTP authentication.
7. Choose the event types to which you would like to subscribe. You may select any number of event types
8. Select **Create**

Managing Webhooks

Once configured, Webhooks can be managed via the **System** tab. On that tab, in the **Advanced** section, select **Webhooks** to see established Webhooks.

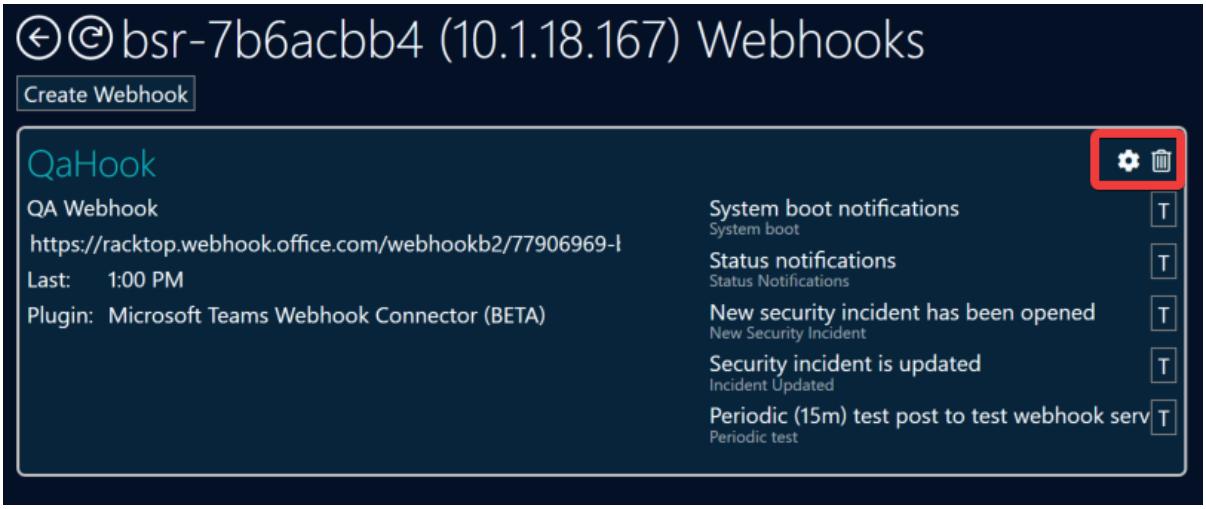

Each Webhook can be reconfigured by selecting the gear icon in the upper left corner. This includes allowing you to make changes to which events are subscribed. The trash icon can be clicked to remove the configured Webhook. To send a test Webhook notification for a specific event, select the **T** icon next to that event. The test notification will be sent immediately to the configured application.

Upgrading

The following topics explain how to upgrade to BrickStor SP Release 23.2 for single-node configurations.

IMPORTANT	Please contact RackTop support for specific upgrade instructions if you are upgrading from release 21.x or prior.
IMPORTANT	Please contact RackTop support to arrange for assistance when upgrading a BrickStor SP HA cluster configuration.
NOTE	If your BrickStor is in an air-gapped network or otherwise unable to download the latest version from the public internet, you can request a download link at https://support.racktopsystems.com or by contacting RackTop support to arrange alternative means.

Pre-upgrade Considerations

There are some important items to note when upgrading from Release 22 to Release 23.2. The email and report system has been greatly improved in Release 23.2, but this means any existing email and report configurations will not be carried forward and will need to be configured after the upgrade is complete. Please review these settings and note them where necessary prior to commencing with the upgrade.

Upgrading a Single Node BrickStor using the latest BrickStor SP Manager

The following steps demonstrate the upgrade process for a single, standalone BrickStor SP configuration.

1. Beginning the upgrade

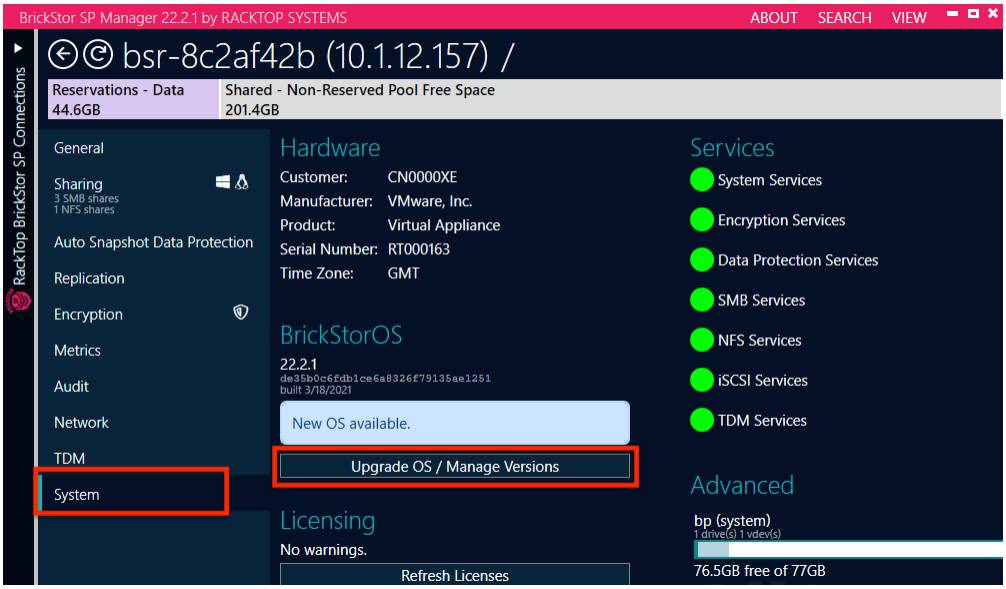

- Connect the BrickStor SP Manager to your appliance.
- Choose **Upgrade OS / Manage Versions** to perform the upgrade.

2. Download the new OS version

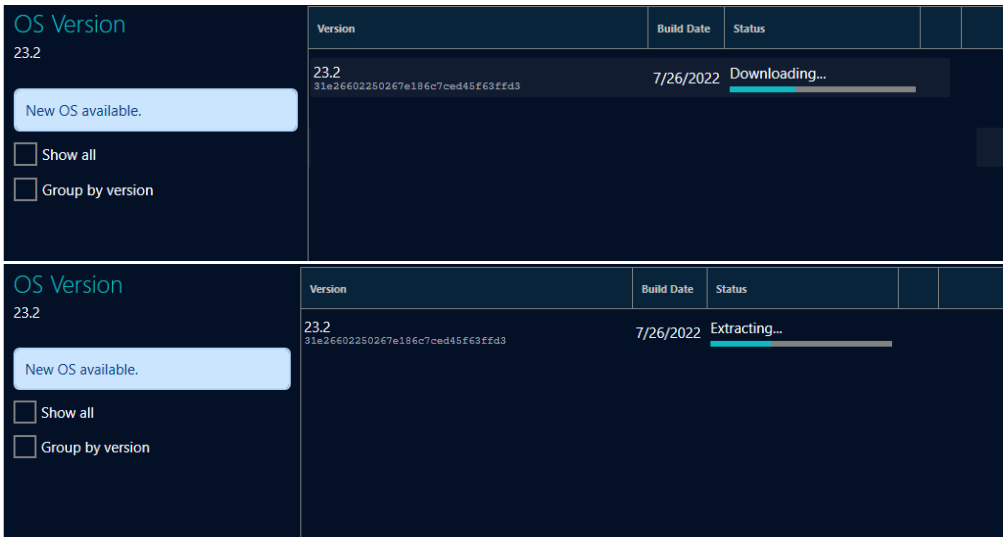

- Choose the version to download by clicking the Download link.

3. Activating OS Version

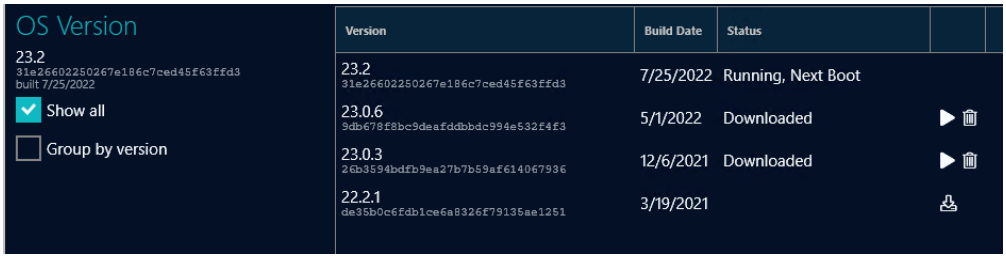

- Once downloaded, click the "play" icon to activate at next boot.

4. Commit the OS Upgrade Change

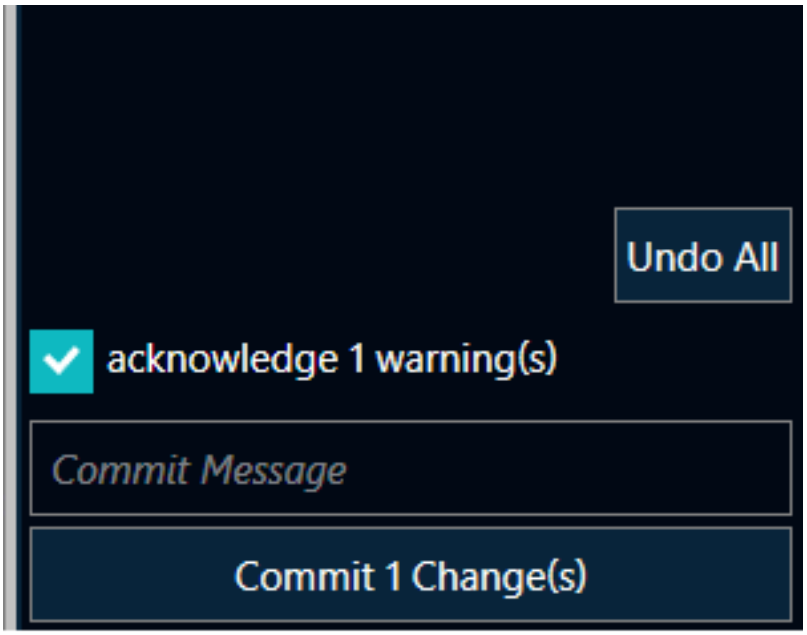

- Commit the change in the Changes pane.

5. Reboot the System

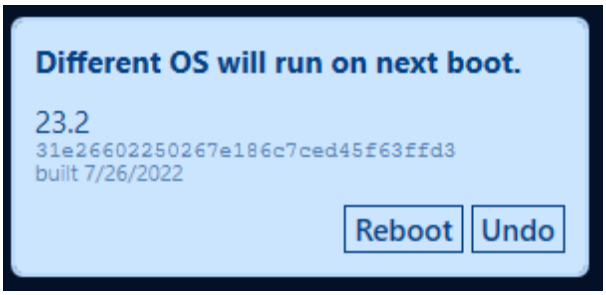

The BrickStor SP appliance will now reboot into the new version of the OS. After it does so, navigate to its IP address or hostname in a web browser and log in. You will be asked to review and accept the Terms & Conditions before proceeding. Once you have done that, you will be able to download the new version of the BrickStor SP Manager.

Post-Upgrade Tasks

Once you are connected to your BrickStor SP system using the new version of the BrickStor SP Manager, be sure to do the following:

- Reconfigure any SMTP email settings.
- Review and configure any desired report settings.
- Review the rest of this documentation for new features that you may wish to configure or activate.

Upgrading Distributed Configuration Database (confd) from 23.0.6 to 23.2

Preparation

Proceed to the following steps to insure prerequisites before the HA and confd upgrade are fulfilled:

- Upgrade both hosts to same version of 23.2 through the UI, balance pools as necessary, and ensure cluster nodes are online.
- Ensure hiavd and confd services are online.
- You have already installed the BrickStor SP GUI 23.2.

NOTE BrickStor SP should prompt two errors, hiavd is running 23.0 instead of 23.2, and a confd issue.

NOTE If not logged onto domain controller, make time source <IP of Time Server> and click **sync now**.

- Make sure that the windows system can ping all members of the cluster, and vice versa.

NOTE Some versions of windows may disable ICMP echo either by default or by policy.

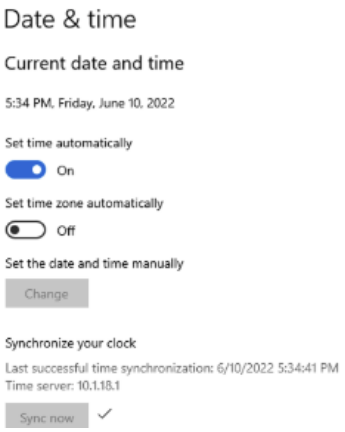

Instance 0

To upgrade a confd from an Instance 0 install on the witness:

- Install confd as Instance 0 - don't start the service

> **NOTE** Instance **must** be specified to 0.

- Disable the existing confd service
- Open the **existing** confd instance folder (eg. `c:\racktop*example*`)
- Copy the entire original data directory to `c:\program files\racktop\brickstor\confd\00\data`
- Copy the .svc.creds from `c:\racktop` to `c:\program files\racktop\brickstor\confd\00\cfg`
- Copy the confd.conf from `c:\racktop` to `c:\program files\racktop\brickstor\confd\00\cfg`
- From the Administrator command line, run `c:\program files\racktop\brickstor\confd\00\confd.exe -svc`
- If you see **authentication errors**, input **ctrl+c**
- From the Administrator command line, run `c:\program files\racktop\brickstor\confd\00\confd.exe -svc -node-acct-recovery`
  - The default user/pass for the etcd database is `root/racktop`.
- From the Administrator command line, run `c:\program files\racktop\brickstor\confd\00\confd.exe -svc`.
- Input **ctrl+c**, and start the system normally.

Addendum

The following section will include various extra information to aid in the setup and use of the BrickStor SP Manager.

Open Network Port Requirements

Table 1. BrickStor SP Open Network Port Requirements

Ports	Description/Service	Protocol	Direction	This port is open to/Purpose
22	SSH	TCP	inbound	Receive Management and Replication data
22, 8444, 8544	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)
443	HTTPS	TCP	inbound	RMM/iLO Out of Band Management
514	syslog	TCP/UDP	outbound	Logging
623	rmcp	TCP/UDP	inbound	HA Power/IPMI access
2049	NFS/portmap	TCP/UDP	inbound	NFS client access
2379,2380	confd	TCP	inbound	Configuration database
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
5697	keymgrd	TCP	bidirectional	Key replication/sync
5699	bsrlicensed	TCP	bidirectional	HA license check
8086, 8088	influxdb	TCP	inbound	Used for BrickStor SP Manager (charts)
80, 443, 8443	bsrapid	TCP	inbound	Used for BrickStor SP Manager (http/https)