Redbooks
ibm.com/redbooks

Policy-Based Replication with IBM Storage FlashSystem, IBM SAN Volume Controller and IBM Storage Virtualize

Chris Bulmer

Chris Canto

Daniel Dent

Bill Passingham

Nolan Rogers

David Seager

Storage

IBM

Redpaper

IBM Redbooks

Policy-Based Replication with IBM Storage FlashSystem, IBM SAN Volume Controller and IBM Storage Virtualize

February 2023

REDP-5704-00

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (February 2023)

This edition applies to IBM Storage Virtualize Version 8.5.2 and later.

This document was created or updated on February 14, 2023.

Contents

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Easy Tier®
FlashCopy®
HyperSwap®
IBM®
IBM FlashCore®
IBM FlashSystem®
IBM Spectrum®
Redbooks®
Redbooks (logo) ®
Storwize®

The following terms are trademarks of other companies:

Ansible, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Preface

Policy-based replication is the successor to Remote Copy for providing replication services for IBM Storage FlashSystem, IBM SAN Volume Controller, and IBM Storage Virtualize for version 8.5.2 and later. This new management model uses volume groups and replication policies to enable the system to automatically deploy and manage replication. This significantly simplifies the tasks that are associated with configuring, managing, and monitoring replication.

Compared to Remote Copy, policy-based replication replicates data between systems with minimal overhead, significantly higher throughput, and reduced latency characteristics.

This IBM Redpaper publication provides a broad understanding of policy-based replication and step-by-step implementation details.

This paper is intended for use by pre-sales and post-sales technical support and storage administrators.

Authors

This paper was produced by a team of specialists from the United Kingdom.

Chris Bulmer is a senior software engineer for IBM Infrastructure based at IBM Hursley in the UK. Chris worked on designing, developing, and testing IBM Storage products since 2013. He led the development of high-availability and replication features for IBM Storage FlashSystem, IBM SAN Volume Controller, and IBM Storage Virtualize since 2017.

Chris Canto is a senior software engineer for IBM Infrastructure based at IBM Hursley in the UK. Chris has almost 20 years of experience in working on IBM Storage products in test, support, and development roles. In his current role, he is part of the team responsible for delivering the next generation of replication and high-availability features, with a focus on usability.

Daniel Dent is a software engineer for IBM Infrastructure based at IBM Hursley in the UK. He has eight years of experience developing and testing IBM Storage Virtualize and worked on a wide range of software features. In his current role, he is the Test Automation lead for policy-based replication.

Bill Passingham is a software engineer for IBM Infrastructure based at IBM Hursley in the UK. He worked with the IBM Storage Virtualize products since their inception in the early 2000s, primarily in a development role, but with experience in field support and testing. He is currently working in the team responsible for replication, specializing in node-to-node messaging technologies.

Nolan Rogers is a software engineer for IBM Infrastructure based at IBM Hursley in the UK. He worked with the IBM Storage Virtualize products since 2008, with a focus on testing a range of features and hardware platforms. He is currently working in the team responsible for replication technologies, most recently policy-based replication.

David Seager is a software engineer for IBM Infrastructure based at IBM Hursley in the UK. He has two years of experience designing, developing, and testing IBM Storage Virtualize and has a total of 25 years of development experience across IBM. In his current role, he is part of the team that is responsible for delivering new replication features. David contributed to the IBM Redbooks publication *CICS Transaction Gateway V5 The WebSphere Connector for CICS*, SG24-6133, and two other IBM® Redpapers.

Thanks to the following people for their contributions to this project:

- Vasfi Gucer
  Austin Center
- Marcela Adan, Wendi Gervis
  IBM Redbooks®

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an IBM Redbooks residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on LinkedIn:

  `http://www.linkedin.com/groups?home=&gid=2130806`

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  `https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm`

- Stay current on recent Redbooks publications with RSS Feeds:

  `http://www.redbooks.ibm.com/rss.html`

1

Introduction

Policy-based replication is the successor to Remote Copy for providing replication services for IBM Storage FlashSystem, IBM SAN Volume Controller, and IBM Storage Virtualize for version 8.5.2 and later. Throughout this document, IBM Storage Virtualize is referred to as the general term for these products as this is the software on which all of these systems run. The policy-based replication management model uses volume groups and replication policies to enable the system to automatically deploy and manage replication. This model significantly simplifies the tasks that are associated with configuring, managing, and monitoring replication. A new ‘ground-up’ implementation is used to support demanding workload applications. This new implementation complements the new management model and replicates data between systems. Compared to Remote Copy, policy-based replication replicates data between systems with minimal overhead, significantly higher throughput, and lower latency characteristics.

Policy-based replication and Remote Copy are technologies that enable you to keep a real-time copy of data at a remote site that contains another IBM Storage Virtualize system. Such a remote site is referred to as the Disaster Recovery or DR system. The site where the data is being accessed by hosts is referred to as the Production system.

Some concepts, such as partnerships, are common between Remote Copy and policy-based replication, whereas other concepts are different. When relevant throughout this document, common concepts and how they are used for both Remote Copy and policy-based replication are noted.

1.1 Asynchronous policy-based replication

The initial release of policy-based replication supports asynchronous replication between two systems. Asynchronous policy-based replication provides a variable, greater-than-zero recovery point that aims to achieve the best possible recovery point for the current conditions. This type of replication ensures mutual consistency between all volumes in the volume group. This is the successor for both Global Mirror and Global Mirror with Change Volumes.

Asynchronous policy-based replication uses two modes of operation that it automatically switches between.

- The first mode acknowledges host-writes as soon as they can be committed to the local storage system, sequence-tagged, and transmitted to the remote system. This technique allows asynchronous policy-based replication to be used over longer distances than synchronous replication. Each asynchronous replication implementation delivers a recovery point greater than zero. However, in this mode, the actual recovery point tends to be small, typically anywhere from several milliseconds to some number of seconds depending on the current conditions.

  This is the default mode and all volume groups that use asynchronous policy-based replication aim to run in this mode whenever possible.

- The second mode is used when insufficient bandwidth exists between systems or throughput at the DR system to support the replication of every write that is received from the application. This mode uses point-in-time snapshots of the production volumes to periodically resynchronize the DR system. Many applications write to only a small portion of a volume at a time, so this approach allows for the coalescing of writes to reduce the bandwidth requirements. However, this mode provides a higher recovery point than the first mode, typically minutes or hours. Therefore, this mode is used only when the first mode cannot be sustained.

2

Planning for replication

This section describes the items to consider when you plan for policy-based replication. If you are planning to use replication that is provided by Remote Copy, see the IBM SAN Volume Controller and Storwize® Family Native IP Replication Redbook at: https://www.redbooks.ibm.com/redpapers/pdfs/redp5103.pdf.

2.1 Business requirements

When you are planning for replication, it is important to understand the business continuity requirements. Typically, these requirements include understanding the following items:

- The Recovery Point Objective (RPO) and Recovery Time Objective (RTO) for the applications.
- The requirements for the number of copies and the geographical separation between systems. Different applications might have different requirements.

The common approach that is used for disaster recovery (DR) is to maintain multiple geographically-separated copies of your data. However, the response times that are experienced by applications often dictate the replication approach that is used. Synchronous replication and high availability must ensure that writes are completed at both systems before completion is acknowledged to the application. This results in a higher round-trip time (RTT) between systems.

The speed of light in glass (in fiber optic cable) has an RTT of approximately 100 KM per millisecond. Therefore, application performance requirements limit the use of synchronous replication to relatively short distances. More delays are typically introduced by switches, routers, and other network devices that must also be considered. It is important to understand the route that the data must take between your sites as the distance of the cables might be significantly higher than the geographical distance and therefore increase the RTT.

In 2003, the US Securities and Exchange Commission (SEC) suggested that 200 miles (320 kms) would be appropriate between the primary and secondary sites for a US financial institution. However, many commercial and government-owned companies are content with distances of 30 km or even less as a compromise between cost, convenience, and protection, because most outages have relatively localized causes.

The following list includes examples of disaster radii that you might want to consider:

- Regional power outage
- Problem that is localized to computer room, building, or urban area
- Typical hurricane 150 km
- 2005 Hurricane Katrina, 250 km sweep
- 2011 Thailand floods, up to 60 km radius
- 2011 Christchurch earthquake, 60 km damage radius
- 2011 Tohoku earthquake and tsunami, 10 km inland
- 2011 Fukushima Daiichi disaster, up to 50 km sweep
- 1980 Mt St. Helens volcanic eruption, up to 30 km radius

The RTT affects the throughput of replication as each replicated write requires resources within the storage array.

- With an RTT of 1 ms, each resource can be used 1000 times per second.
- With an RTT of 250 ms, each resource can be used only four times per second.

Storage arrays have only a finite number of resources available for replication. Therefore, the greater the RTT, the less the maximum throughput.

Another aspect to include in planning is the change rate of the data that you need to protect. High change rates with low-bandwidth result in the system not being able to achieve low RPOs.

Budgets must also be considered as part of planning for business continuity as the cost that is associated with network bandwidth can vary between countries and regions. Ultimately, cost is a deciding factor when you are planning for business continuity. This can be the cost of deploying and maintaining DR or the business cost of downtime if a disaster occurs. To help with the planning, you must have a good understanding the budget, the geographical separation requirements, and the required RPO.

2.2 Disaster recovery

Disaster recovery that uses asynchronous replication allows for much greater geographic separation between the systems. This type of DR has a greater-than-zero RPO and RTO as it disconnects the completion of the production write with replicating the data to the DR system. It is also capable of replicating by using lower-bandwidth links between the systems, but it achieves a higher recovery point because the network acts as a bottleneck.

If you require either zero RPO (no data loss) or zero RTO (no downtime), then you need synchronous replication or high availability. At the time of writing, high availability with IBM Storage Virtualize is provided by either IBM HyperSwap® or Stretched Cluster and synchronous replication is provided by Metro Mirror. These options are covered in the IBM Redbooks publication Implementation Guide for IBM Spectrum Virtualize Version 8.5, SG24-8520-00.

Disaster recovery that uses replication protects against large-scale disasters that affect a site or storage array and provides a write-order consistent set of data. This type of DR does not protect against corruption or cyberattack. IBM FlashCopy® can be used to maintain point-in-time copies to protect against corruption. IBM Safeguarded Copy can be used to protect production data against cyberattack while it provides quick recovery from a previous copy.

2.3 Planning for applications

When you are planning to set up replication, consider the items that are described in the following sections. Replication can be managed by using external tools, such as Ansible, that can be integrated into workflows.

2.3.1 Multipath configuration

IBM Storage Virtualize systems that are active/active as both nodes in an I/O group are capable of processing host I/O. These systems use SCSI ALUA (Asymmetric Logical Unit Access) or NVMe ANA (Asymmetric Namespace Access) to load balance and improve read cache hit efficiency. The preferred node has the responsibility for performing replication operations for a volume. When a write is received from a host, the preferred node allocates the resources for replication, performs metadata updates, and transmits the data to the remote system. By default, the preferred node is selected automatically by the system when a volume is created to load-balance between the nodes in the I/O group. However, the preferred node can be selected manually and can be changed later if required.

If a host write is received on a non-preferred node for a volume that uses policy-based replication, the host write is forwarded internally to the preferred node for processing. When all replication processing for a volume is completed by the preferred node, the benefits include:

- Optimizations that reduce the write latency that is experienced by applications.
- Best use of the available resources within the storage system.

If the preferred node is unavailable, such as during maintenance activities, the processing is automatically transitioned to the partner node in the I/O group. Processing automatically transitions back when the preferred node returns.

For these reasons, it is recommended that you use SCSI ALUA or NVMe ANA for hosts that use policy-based replication. If you are using SAN Volume Controller enhanced stretched cluster, it is recommended that you change the preferred node to the site that generates most of the host writes in normal running.

> **Note:** Most operating systems default to using ALUA-compliant or ANA-compliant multipath policies.

2.3.2 Crash consistency

When you copy the contents of one volume to another, it must contain a write-order consistent set of data for the copy to be of value. This is sometimes referred to as crash-consistent. That is, the data set is in the same state it would be in if the power failed and the system came to a sudden unplanned stop. File systems and databases are designed to be able to cope with their logical disks being in such a state and run recovery routines such as `fsck` (file system consistency check) and replay redo logs.

IBM Storage Virtualize replication is designed to always maintain write-order consistency in all operations, except for synchronization. When volumes are synchronized, either initially or after a disruption, the synchronization catch-up process is carried out without regard to the original write order. The consistency of the recovery copy is protected automatically by using the associated change volume during the synchronization to ensure that the DR volumes are always available to be accessed if required.

Policy-based replication uses volume groups to define the set of volumes that are required to be mutually consistent. Typically, a volume group should contain all the volumes for an application that needs to be recovered in a disaster.

2.3.3 Operating system support

Asynchronous policy-based replication is supported on any operating system and protocol that is supported by IBM Storage Virtualize. The full interoperability list can be found at `http://www.ibm.com/systems/support/storage/ssic`.

2.4 Partnerships and connectivity

Partnerships are used to connect two IBM Storage Virtualize systems together so that they can replicate and perform configuration tasks on each other. Each system supports up to three partnerships to other systems connected over Ethernet or Fibre Channel.

2.4.1 Fibre channel partnerships

Additionally, a new requirement for Fibre Channel partnerships that is used for policy-based replication is for the partnered systems to have IP connectivity between the management IP addresses of the systems. This IP connectivity uses the Rest API that requires inbound connections on port 7443. The replication data uses the Fibre Channel infrastructure, but a small amount of management traffic requires the IP connection to perform configuration activities. The management traffic uses the RESTful API; it requires certificates to be configured for authentication and uses SSL to encrypt traffic between the systems.

FCIP routers, Dense Wave Division Multiplexing, or Complex Wave Division Multiplexing can be used for the inter-system link.

Zoning requirements

Two supported zoning configurations exist for partnerships that use policy-based replication, depending on the round-trip time (RTT) between the systems.

- RTT between sites is <= 80 ms: Use zoning and port masking (optional) such that nodes within I/O groups that are specified by a replication policy can communicate directly with all nodes in I/O groups that are the target of those replicated volumes in the partnered system.
- Round-trip latency between sites is >80 ms and <= 250 ms: Zoning should be applied to provide separate intersystem zones for each local-remote I/O group pair that is used for replication.

Any zoning configuration that requires replication I/O to be forwarded between local nodes or remote nodes is not supported, but can be tolerated during failures or maintenance.

The objective of the extra zoning constraints for the higher round-trip times is to remove connections that would otherwise be idle. These connections can consume intersystem resources that can be used more efficiently elsewhere if they are not present.

Figure 2-1 shows a diagram of recommended zoning. This example contains two replication policies:

- First replication policy: Specifies I/O group 0 in system 1 and I/O group 0 in system 2.
- Second replication policy: Specifies I/O group 1 in system 1 and I/O group 1 in system 2.

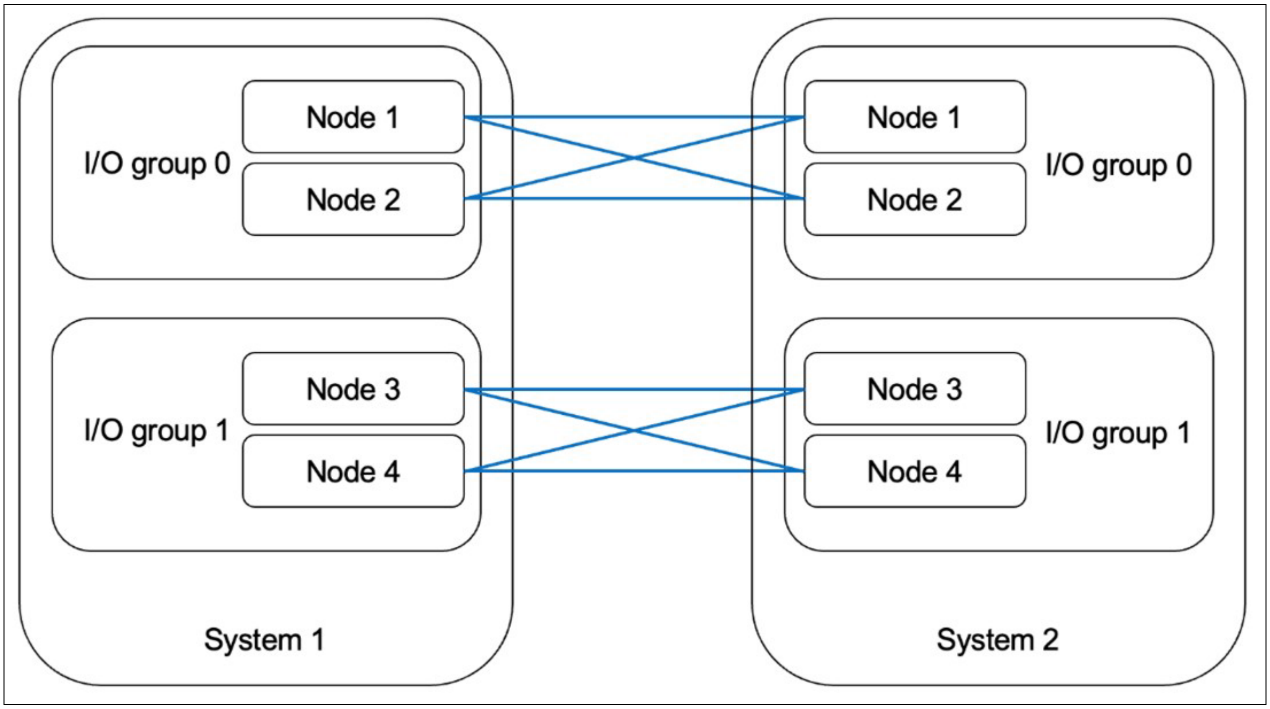

Figure 2-1 Additional zoning required when adding replication to another I/O group

If another replication policy exists that specifies I/O group 0 in system 1 and I/O group 1 in system 2, then connectivity is required between both nodes in I/O group 0 in system 1 and I/O group 1 in system 2. The additional required zoning is shown in Figure 2-2:

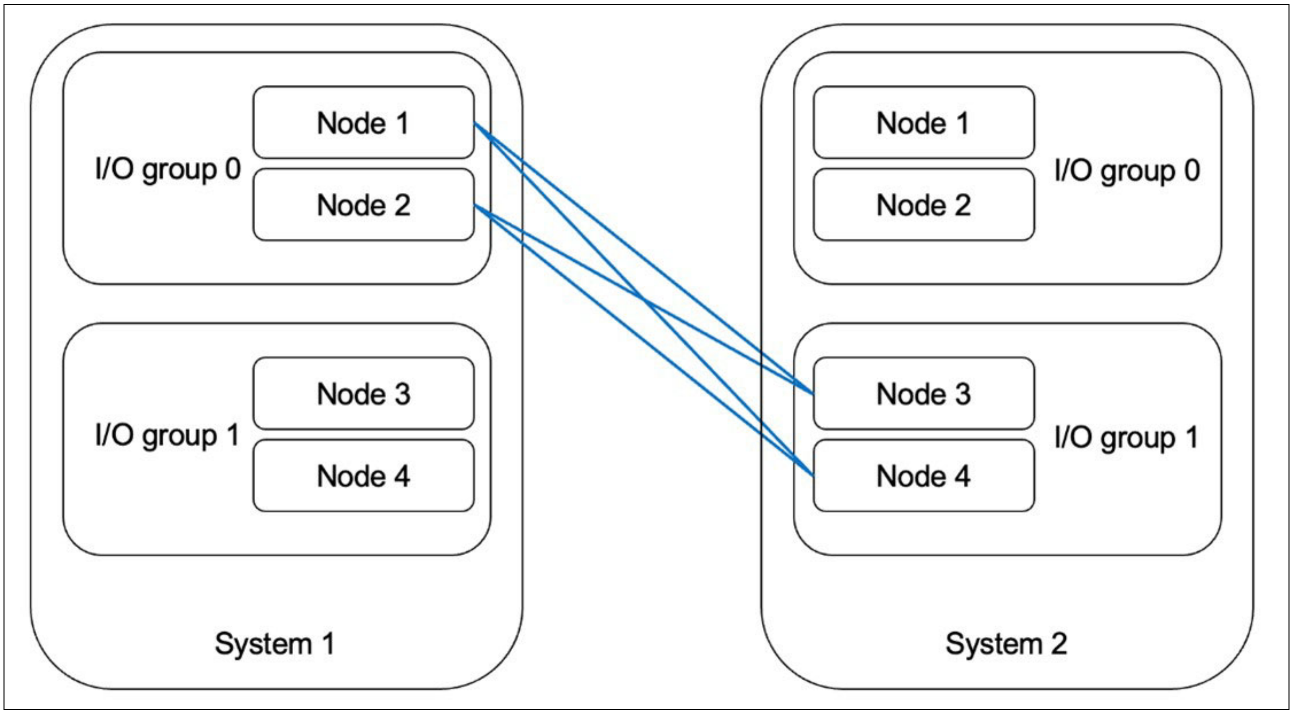

Figure 2-2 Additional zoning required when adding replication to another I/O group

Fibre Channel port assignment

In all cases, either two or four Fibre Channel ports from each node should be used for replication and these ports should not be used for any other purpose. The only exception is on systems with four Fibre Channel ports per node. In this case two ports might be shared with local system traffic.

A combination of zoning and system-level Fibre Channel port masking can be used to meet the port-assignment requirements.

For more information of Fibre Channel port assignments, see *Performance and Best Practices Guide for IBM Spectrum Virtualize 8.5* at:
https://www.redbooks.ibm.com/abstracts/sg248521.html.

2.4.2 IP partnerships

Native IP partnerships are supported by policy-based replication and support a maximum RTT of 80 milliseconds. The minimum bandwidth requirement for the inter-site link is 10 Mbps. While the system operates in the presence of packet loss, increasing amounts of packet loss affect the rate at which data can be replicated.

Transmission Control Protocol (TCP) ports 3260 and 3265 are used by systems for IP partnership communications. Therefore, these ports need to be open.

- Port 3260 is used for link initialization and system discovery.
- Port 3265 is used for the IP replication traffic.

Connectivity must be permitted between the cluster management IP addresses in each system and in the ports that are selected for use for replication.

More details for IP partnership requirements are in the documentation, including configuring ports for replication and portsets.

The system can optionally compress and encrypt the data in flight when it uses IP partnerships.

2.4.3 Layers

Layering is a concept that protects against erroneous configuration between IBM Storage Virtualize systems. Two layers exist: Replication and Storage. Partnerships can be created only between systems that have the same layer (both are in the replication layer or both are in the storage layer). Layering is also important for external virtualization as only systems in the storage layer can present volumes to an IBM Storage Virtualize system in the replication layer.

Note: A system is not required to be in the replication layer to use replication. The requirement is that both systems are in the same layer.

IBM Storage FlashSystem arrays default to the storage layer and IBM SAN Volume Controller systems default to the replication layer. Changing the layer of the system is generally performed only during initial setup of the system. Most systems never change the layer from the default setting.

3

Introducing policy-based replication

This section describes the concepts and design of policy-based replication. It provides the implementation details of asynchronous replication and useful information for understanding how it functions.

3.1 Configuration

One of the key goals of the new configuration model is that external automation can perform provisioning activity on the local system as they would for a non-replicated solution. The system is responsible for performing all the steps that are required to enable the replication function on the volumes. This approach also significantly improves the experience for interactive users who use the CLI, GUI, or REST API as it reduces the number of repetitive steps that are required to setup replication. This approach also ensures that replication is configured consistently and as intended.

To achieve this new configuration model, the system must capture the replication requirements from the storage architect ahead of time, in the form of policies.

Configuration of replication requires five parts:

- **Partnerships** for allowing access between a pair of systems.
- **Replication policies** for defining replication settings. The policies are assigned to volume groups.
- **Volume groups** for managing replication of groups of volumes.
- **Provisioning policies** for defining how volumes should be provisioned within a pool. the policies are assigned to storage pools.
- **Storage pool links** for defining where capacity should be provisioned from on remote systems.

3.1.1 Partnerships

A partnership allows for configuration and data to be exchanged between two systems for the purposes of replication.

When policy-based replication is used, IP connectivity is required between the management IP addresses of each system. A new requirement for partnerships that are used for policy-based replication is the installation of Transport Layer Security (TLS) certificates between partnered systems. The certificates are required to provide secure communication between systems for managing replication between the systems. Management traffic is authenticated by a certificate and is encrypted using SSL. It is recommended that you create partnerships by using the GUI as this automates the certificate exchange and setup. Both self-signed and signed certificates are supported.

The maximum amount of bandwidth that is consumed by policy-based replication is controlled by the settings on the partnership. Differences between policy-based replication and Remote Copy exist in how the partnership settings are used.

The partnership bandwidth defines the amount of bandwidth that is dedicated for replication between each of the I/O groups in megabits per second.

The background copy rate defines the amount of the bandwidth, expressed as a percentage that can be used for synchronization. For example, a link bandwidth of 1000 Mbps and a background copy rate of 50 allows for 500 Mbps of synchronization traffic (62.5 MBps). This control is a separate parameter because when Remote Copy is used, this control must be managed by the user to balance foreground host writes and background synchronization activity. Policy-based replication does not distinguish between the two cases and automatically manages the bandwidth for each type.

If you are using only policy-based replication, then the background copy rate should be set to 100% to allow all available bandwidth to be used.

If you are using both Remote Copy and Policy-based replication on the same partnership, then the background copy rate can be used to balance the bandwidth between policy-based replication and Remote Copy.

All traffic for policy-based replication is treated as background copy, so the rate can be managed to guarantee bandwidth to Remote Copy, which is more sensitive to changes in the available bandwidth.

Remote Copy has additional system settings for the following items:

- To control the synchronization rate of an individual relationship.

  This setting is not used for policy-based replication; the system automatically manages the synchronization bandwidth.
- To control the amount of memory used for replication.

  This setting is not used for policy-based replication.

3.1.2 Replication policies

Replication policies are a key concept for policy-based replication as they define how replication should be configured between partnered systems. A replication policy specifies how replication should be applied to a volume group and therefore to all volumes within that group. A replication policy can be associated with any number of volume groups. A volume group can have at most one replication policy associated with it.

A replication policy defines three key attributes:

- **A set of locations** defining which I/O groups on which of the partnered systems should contain a copy of the volume group, where the data should be replicated.
- **A topology** representing how the systems are organized and the type of replication that should be performed between each location, how the data should be replicated between the locations.
- **A name** to uniquely identify the replication policy on both systems.

Replication policies are immutable and therefore cannot be changed once they are created. This is important as it guarantees that both systems always have the same definition of how replication should be configured. If changes are required to a replication policy, a new policy can be created with the desired changes, and the associated volume groups can be reassociated with the new policy. Each system supports up to 32 replication policies to be created, which allows for various locations, topologies, and RPOs to be defined.

The creation of a replication policy results in the system creating the policy on all systems that are defined in the replication policy. A replication policy can only be created when the systems are connected. A replication policy can be deleted only if there are no volume groups that are associated with the policy and can be deleted while the systems are disconnected.

2-site-async-dr topology

The initial release supports one topology: `2-site-async-dr`. This topology defines that two locations are involved (2-site) and asynchronous replication is being performed for disaster recovery.

For this policy an additional attribute is required: an RPO alert. The system frequently checks the current recovery point of each volume group and if the current value exceeds the objective, it triggers an event to be raised. This allows you to easily monitor replication and be aware if it is outside of your defined RPO. The RPO value in the replication policy is also used internally to prioritize the mode in which volume groups should operate if bandwidth between systems is constrained.

3.1.3 Volumes and volume groups

Volume groups provide a method for grouping volumes that are used by an application. Replication is configured on the volume group by assigning a replication policy to the volume group. The system then automatically performs actions to configure replication according to the replication policy for all volumes in the volume group. This simplifies the configuration of replication as it is only configured once. Adding a volume to a group automatically configures replication for it. It also permits configuration changes to be performed while the partnership is disconnected. The system automatically reconfigures the DR system once the partnership is reconnected.

Volume groups can also be used without replication or with other copy services such as FlashCopy, Safeguarded Copy, or Transparent Cloud Tiering. The same volume group might be used for multiple features (for example, replication and Safeguarded Copy), which allows for a single definition of the volumes that are required by an application. Restrictions are listed in the IBM FlashSystem 9x00 documentation at:
https://www.ibm.com/docs/en/flashsystem-9x00.

> **Note:** This document contains links to IBM Documentation. Our starting point is the IBM Storage FlashSystem 9x00 family web page. However, the reader might need to select the product that applies to their environment. See “Related publications” on page 63 for a list of the Product Documentation links for IBM Storage Virtualize.

> **Note:** Associating a replication policy to a volume group before you create new volumes allows the initial synchronization to be skipped. Therefore, if you are intending to provision new volumes, it is recommended that you assign a replication policy to the volume group first, then create the volumes in the volume group.

Separate relationships or consistency groups are not used for policy-based replication. Configuration information is presented on the volume or volume group and is managed by using the volume group.

Three replication modes exist for volume groups:

- **Production**: Volumes in the volume group are accessible for host I/O and configuration changes are allowed. This copy acts as the source for replication to any recovery copies. By definition, a production volume group always contains an up-to-date copy of application data for host access.
- **Recovery**: Volumes in the group are able to receive only replication I/O; volumes act as a target for replication and cannot be used for host I/O. Most configuration changes are not permitted and must be performed on the production copy.
- **Independent**: If independent access is enabled for the volume group, such as during a disaster, each copy of the volume group is accessible for host I/O and permits configuration changes. Replication is suspended in this state.

Replication policies do not define the direction in which the replication is performed. Instead, the direction is determined when a replication policy is associated with a volume group, or a volume group is created by specifying a replication policy. The system where this action was performed is configured as the production copy of the volume group.

To change the direction of replication, the volume group must first be made independent. After the volume group is independent, the direction can be selected by choosing which location should be the production copy of the volume group and then restarting replication from that system. If you wish to perform a planned failover, it is important to suspend application I/O and ensure that the recovery point is more recent than when the application stopped performing I/O. Failure to do this can result in data missing after the change of direction.

An important concept is that the configuration and the data on the volumes are coupled and the recovery point is formed from both the configuration and volume data. Adding, creating, removing, or deleting volumes from a volume group occurs on the recovery system at the equivalent point-in-time as they did on the production system. This means that if independent access is enabled on a volume group during a disaster, then the data and volumes that are presented are as they were from a previous point in time on the production system. This might result in partially synchronized volumes being removed from the recovery system when independent access is enabled, if those volumes are not in the recovery point.

When a volume is deleted from the production volume group, the copy of the volume from the recovery volume group will not be immediately deleted. This volume is deleted when the recovery volume group has a recovery point that does not include that volume.

This coupling requires that all configuration changes to the volume group are made from the system that has the production copy of the volume group and the changes are reflected to the recovery system. The exception to this requirement is the action of enabling independent access on a recovery copy of a volume group as this action can be performed only on the system that includes the recovery copy of the volume group.

When a replication policy is removed from a volume group, the recovery volume group and its volumes are deleted. If you need to keep the recovery copy, independent access should be enabled on the recovery copy first. Then, the replication policy can be removed from the volume group.

The replication policy that is associated with a volume group might be changed to a compatible policy without requiring a full resynchronization. A compatible policy is defined as having the same locations and allows for a replication policy to be replaced by another policy with a different recovery-point objective without interrupting replication. If you attempt to change the policy to an incompatible policy, it is rejected by the system. The replication policy needs to be removed and the new policy must be associated with the volume group.

3.1.4 Pool linking

Storage-pool linking provides a mechanism to define which storage pool the system should create copies of a volume in on remote systems. A pool can be linked to only one pool per partnership. A pool link is required to be configured for a partnership to create replicated volumes in that pool. The pools can be parent pools or child pools.

For a standard topology system, if a volume is created locally as a mirrored volume between different pools, the system uses the pool link of the primary volume copy to identify the remote pool to use. If the system topology is stretched, then the pool in site 1 is used to identify the linked pool; pools in site 2 are not required to be linked to replicate stretched volumes. The system does not provide a way to automatically create mirrored volumes on a remote system, but mirror copies can be added by the user.

3.1.5 Provisioning policies

Replication policies define the remote system and the I/O group in which the system should create the recovery copy of the volume group. Provisioning policies complement this by allowing storage architects to define how the new volume capacity should be provisioned within a storage pool.

When combined with storage-pool linking, the system understands the following items:

- Which storage pool in which system and I/O group a volume should be created.
- Which attributes to use to create the volume.

  These attributes define the additional capacity savings to be applied to the volume and any built-in capacity savings provided by the storage array. For example, pools containing FlashCore Module (FCM) drives are self-compressing. Additional options might be added to provisioning policies in the future.

A provisioning policy can be associated with multiple pools. A pool can be associated with, at most, one provisioning policy. Provisioning policies are supported on parent pools and child pools. A provisioning policy exists within a single system (unlike replication policies, which are replicated between systems). Intentionally, few configurable options exist in the provisioning policy. The defaults for the other attributes are chosen to match the most common usage and best practices.

The association of the provisioning policy to a pool automatically causes it to be used for new volumes that are created in the pool. This means that the provisioning policy is used simply by specifying a pool that has a policy that is associated when creating a volume by using the GUI or the `mkvolume` command. Provisioning policies can also be used for non-replicated volumes to simplify volume creation.

Two provisioning policies are created automatically when the first parent pool is created.

- One pool is optimized for performance and creates fully-allocated volumes.
- One pool is optimized for capacity and employs all the data reduction features of the storage pool.

If IBM FlashCore® modules are used in the storage pool, then fully-allocated volumes still benefit from compression performed by the FlashCore module.

Note: Provisioning policies are not used when change volumes are created for policy-based replication as these are always created by the system by using best practices.

3.1.6 Child pools

A single RAID array or mdisk is allowed to be added to only one storage pool. Therefore, it can be linked to only one pool per system and can have only a single provisioning policy associated. Fortunately, child pools can be used to allow for different combinations of pool links and provisioning policies that use the same RAID array or mdisks.

A typical deployment for IBM Storage FlashSystem is to have a single Distributed RAID 6 array with a single storage pool. This parent pool might have multiple child pools created:

- A child pool with a provisioning policy for fully allocated volumes
- A second child pool for compressed volumes
- A third child pool for compressed and deduplicated volumes
- Additional child pools for Safeguarded Copy or multi-tenancy, if required

Each of the pools that contain production volumes must be linked to an equivalent pool in the recovery system.

3.2 Asynchronous replication

If a replication policy is created with the *2-site-async-dr* topology, the system configures asynchronous replication between the systems and I/O groups that are defined in the replication policy. This section describes the design of the asynchronous-replication function that is used by policy-based replication. The design of Remote Copy Global Mirror and Global Mirror with Change Volumes is similar to that of asynchronous replication. However, the implementation differs with the addition of improvements, optimizations, and simplifications.

Like Remote Copy, one system acts as the production location for a volume and is host accessible. The other system acts as a recovery location and is not accessible to the host. Unlike Remote Copy, locations are referred to by their system name rather than *master* and *auxiliary* labels.

3.2.1 Production overview

The production system has two main responsibilities: processing host I/O requests and replicating the contents of the volume to the remote system. The production system can be identified by the location that is shown as production for the volume group.

The production copy can operate in the following three modes:

- **Change recording mode** where new host writes are tracked, written locally but are not replicated.
- **Journaling mode** where new host writes are tracked, written locally, and replicated in-order. This mode has a lower recovery point than cycling mode, but requires more bandwidth.
- **Cycling mode** where new host writes are tracked, written locally and periodically replicated from a snapshot of the production volume recorded on a change volume.

Change recording mode is the default mode and is selected whenever the volume group cannot replicate; this could be because of offline volumes, errors, pending configuration changes, or if the partnership is disconnected. In this mode, the system records which regions of the volume have been written to and need to be synchronized when replication resumes.

When replication becomes possible, the production copy automatically selects either *journaling* or *cycling mode* to perform replication based on which is best for the current conditions. This auto-tuning allows for the system to dynamically switch between the two modes, attempting to achieve the lowest possible recovery point while also avoiding latency problems for the production volumes. For more information on the modes and switching between them, see the following sections:

- “Journaling mode” on page 21
- “Cycling mode” on page 22
- “Mode switching” on page 22

Write tracking and overlapping writes

Each replicated volume is allocated non-volatile memory to track unsynchronized volume regions, known as a bitmap. Both nodes in an I/O group include a copy of the bitmap and updates to the bitmap are mirrored between the nodes. Every volume is split into regions that are called *grains*; each grain is a contiguous 128 KiB range and each bit in the bitmap represents one grain.

Whenever a write is received from an application, the grain is recorded as unsynchronized. The bitmap is set to `synchronized` for a grain once the recovery volume contains identical data. The bitmap is primarily used for to identify regions of the volume that need to be synchronized if replication is suspended. The bitmap also allows for overlapping writes to be detected and resolved.

If one or more writes to the same region of the volume are active at the same time, it is known as an *overlapping write*. It might sometimes be referred to as a *colliding write*. The storage system needs to guarantee that the writes are applied in the same order on both systems to maintain data consistency. In theory there should never be two active writes to the same logical block; this behavior is described as invalid by storage specifications, but in practice it can happen and must be handled by the storage array. To guarantee the ordering of overlapping writes, the system selects one of the writes to process first. This write must be completed before the processing of the second write is started. For best possible performance, application workloads should be aligned to 8 KiB boundary or multiples thereof (16 KiB, 32 KiB, and so on).

Sequencing

In journaling mode, every write is tagged with a sequence number that defines the order in which the writes must be replayed at the recovery system to maintain consistency. Sequencing is across all volumes in a volume group to maintain mutual consistency. To achieve this, one node in the production I/O group is selected to generate sequence numbers for the volume group. As a write is written into the write cache, a parallel process is requests a sequence number for the write. Once the sequence number is obtained and the local write is complete, the host write is eligible for replication.

For performance reasons, sequencing is performed only within an I/O group and all volumes in a volume group that uses policy-based replication must be in the same I/O group. A common cause of performance problems with Remote Copy Global Mirror was the requesting sequence numbers as the sequence number generator could be any node in the system. In this new model, it is either the node that received the write or its I/O group partner node that generates the sequence numbers. The system has significantly better performance characteristics when it sends messages to its I/O group partner node than any other node in the system, which results in better performance and a more consistent solution. Each volume group operates independently and there is no coordination of sequence numbers between volume groups.

Synchronization operations are also tagged with a sequence number, so they are correctly interleaved into the stream of writes.

Journal

Each node maintains a large, volatile memory buffer for journaling host writes and synchronization reads. The exact amount of journal capacity varies between models but (at the time of writing) it can be between 1 GiB and 32 GiB per node. The journal size is related to the number of CPU cores and memory in the node.

The purpose of the journal is twofold.

- Providing temporary in-memory storage for host writes or synchronization reads until they can be replicated to the recovery copy and written to the local cache.
- Buffering host writes in journaling mode so that replication and the local host write are decoupled, which means replication problems cause the recovery point to extend instead of delaying I/O for the production application.

Writes are replicated from the journal to the recovery system in sequence number order. On a per-recovery system I/O group basis, writes are replicated in a first-in-first-out basis to ensure that the recovery system can always make progress. Writes are replicated to the preferred node of the volume in the remote system. If a direct connection is not present between the two nodes, it is routed through local or remote nodes. If the preferred node is not available in the remote system, the write is replicated to the online node in the caching I/O group for the recovery volume.

To ensure that replication does not cause application performance problems, the journal usage is constantly monitored by the system so it can proactively ensure that there is free capacity in the journal for new host writes. The journal is divided by CPU core and maintains lists of writes per volume group. However, resources are not divided between volume groups as this could lead to unfairly penalizing volume groups with higher write workloads.

If the throughput of replication is less than the throughput of the local writes, then the journal starts to fill up. The journal size is determined so that it can accommodate temporary bursts in write throughput. However, if it is a sustained increase in the write throughput versus what can be replicated, a proactive journal purge is triggered to prevent exhausting the resources. All volume groups that are replicating in the I/O group have their journals evaluated, within the context of the RPO defined on the replication policy, for peak and average usage and evaluated. Volume groups that are associated with higher RPO replication policies and those consuming the most journal resources are contenders to be purged first.

Purging a journal involves discarding the journal resources recording host writes that have completed locally, but not yet replicated for the volume group. The volume group then uses the bitmap to synchronize any grains that were affected by writes that were in the journal.

Volume groups have their journals purged if the volume group encounters an error that prevents replication, such as offline volumes or changes in connectivity between the two systems.

3.2.2 Recovery overview

The recovery system runs a process that replays the writes in the order that they were processed at the production location and maintains mutual consistency for the volumes in the volume group.

As writes are received from the production location, the recovery system orders them based on their sequence number. The writes are then mirrored between the two nodes in the recovery system I/O group and written to the volume once all earlier writes are mirrored and processed. Non-overlapping writes can be submitted in parallel to optimize performance but overlapping writes are replayed in sequence number order to ensure consistency. Only one node submits a write, usually the preferred node, but both nodes track its progress through the process.

Writes are recorded in non-volatile memory on both nodes in the I/O group so that they can be recovered if the nodes restart. After the write is stored in non-volatile memory on both nodes, it is marked as committed. Committing a write guarantees that it will be written in the future to the volume and that it forms part of the recovery point.

This is a simplified overview of the process that provides a high throughput method of replaying the writes, which is always able to establish a mutually consistent recovery point for all volumes in a volume group.

3.2.3 Change volumes

Every replicated volume has a change volume associated; these are managed automatically by the system and are hidden from the CLI and GUI views as there are limited user-interactions permitted to them.

A change volume is a thin-provisioned volume or a thin-provisioned and compressed volume if the volume it is associated with is in a data-reduction pool. If all copies of the volume are deduplicated then the change volume is also created as deduplicated, but not available as deduplication source. Change volumes are in the same I/O group with the same preferred node as the host-accessible volume and are created in the same pool; if a mirrored volume is created, then a copy of the change volume is created in both pools. They are always created as cache-enabled volumes with default Easy Tier® settings.

Two FlashCopy maps are created between the volume and the change volume so that these maps can be used to either take or restore a snapshot of the volume. Change volumes do not require or contribute towards the FlashCopy license.

It is crucial to consider the capacity of change volumes when provisioning a system that uses replication. The used capacity of change volumes changes significantly depending on the amount of data that is required to preserve the snapshot. In most cases the data consumes a small percentage of the virtual capacity of the volume. However, in extreme cases it is possible for the data to consume an equal capacity to the volume. A general rule is to size the system with an additional 10-20% capacity of the replicated volumes for change volumes. However, this can vary by system. If there is low bandwidth between sites, but a high write-throughput at the production location, then anticipate the need for more capacity for change volumes. Similarly, if you are planning for the systems to be disconnected for an extended time, the synchronization when they reconnect might need to copy a significant amount of data. The change-volume capacity grows according to the amount of synchronization that is needed.

For example, if a system has 25 volumes that are each 40 GiB, the virtual capacity of that system is 1000 GiB. The amount of usable capacity that is used depends on the data that is written and the data reduction features that are used. Twenty-five change volumes are created in each system that, by default, consume negligible usable capacity. If data reduction features are not used, it would be best practice to have between 1.1 and 1.2 TiB of usable capacity in this example.

Note: If a change volume runs out of space, it might prevent replication.

The FlashCopy maps that are used by the change volumes are automatically started and stopped, as required, to achieve the desired behavior for replication.

In general, the change volume is automatically maintained by the system in line with the volume it protects. However, special cases exist for mirroring and migration when you use use the CLI.

If a volume is migrated to a different pool, the change volume must also be migrated. When you use the GUI, this migration happens automatically. However, a CLI user is required to migrate the change volume in the same way as the user volume. The change-volume vdisk ID can be seen in the **`lsvdisk`** view of the user volume in the **`changevolume1_id`** field.

Note: A CLI user can also specify the **-showhidden** parameter on the CLI views to display the hidden change volume vdisks and FlashCopy maps.

Similarly, for mirroring operations that add or remove a volume copy, the GUI automatically applies the same operation to the change volume, system, which keeps it aligned with the user volume. A CLI user must use the **addvdiskcopy** and **rmvdiskcopy** commands to manage mirrored copies of the change volume when the user volume is modified. When the **addvdiskcopy** command is used on a change volume, the only optional parameters that can be specified are those relating to the mirroring process itself (for example, **autodelete** and **mirrorwritepriority**). This ensures that the change volume is created according to best practice. All other parameters for the new copy of the change volume are defined by the system.

3.2.4 Synchronization

The system automatically performs a synchronization whenever the volume group transitions from change recording mode to either journaling or cycling mode. Before any synchronization starts, the FlashCopy maps for the change volumes at the recovery system are started to preserve the consistency of the volume group. The recovery point is frozen at the time of the last write before the synchronization during the synchronization. This is required because synchronization does not replay writes in order. Therefore, it does not maintain consistency while it is in progress.

Synchronization uses the bitmap to identify the grains that need to be read from the production volumes and written to the recovery volumes. Once the change volumes are maintaining a snapshot of the recovery volumes, the synchronization process starts with the recovery copy controlling the requests for grains to read from the production copy. Synchronization reads are interleaved with host writes (if in journaling mode) and added to the journal. The production system automatically manages the amount of synchronization based on the usage of resources available in the journal. This ensures that synchronization activity does not consume too many journal resources, thus avoiding journal purges that are caused by synchronization.

If synchronization encounters a read error, such as a medium error, the volume group stops replicating and an error is logged against the volume group. After the problem that caused the read error is resolved, replication restarts when the error is marked as fixed. If synchronization is interrupted, such as by the partnership disconnecting, it restarts automatically when it is able and the change volume continues to protect the original snapshot. The snapshot for the change volume is discarded only after all volumes in the volume group complete synchronization or the recovery copy of the volume group is deleted.

If the recovery copy of a volume group is made independent during a synchronization, the change volume is used to restore the snapshot onto the host-accessible volumes. Once the FlashCopy maps are reversed, the volumes are made accessible to the host. This triggers a background process to undo any writes to the volumes that were done as part of the partial synchronization.

3.2.5 Journaling mode

Journaling mode provides a high-throughput, low recovery point asynchronous replication mode. Typically, the recovery point is expected to be below one second, but it can vary based on the RTT, the available bandwidth between systems, and the performance of the recovery system.

In this mode, every host write is tagged with a sequence number and added to the journal for replicating. This is similar in characteristics to Remote Copy Global Mirror, but journaling mode has a significantly greater throughput and the ability to extend the RPO to avoid performance problems.

3.2.6 Cycling mode

Cycling mode uses synchronization and change volumes to periodically replicate to the recovery system in a way that requires less bandwidth and consumes fewer journal resources. Periodically, a snapshot is captured (using FlashCopy) of all volumes in the volume group. The snapshot is stored on the change volume at the production system and is maintained by using either Copy-on-Write or Redirect-on-Write, depending on the storage-pool type and volume-capacity savings. The background synchronization process copies only the changes to the remote system. After the synchronization is complete, the snapshot at the production system is discarded and the process might repeat.

Cycling mode results in a higher recovery point than journaling mode, but has the distinct advantage that it can coalesce writes to the same region of the volume, thus reducing the bandwidth required between systems.

3.2.7 Mode switching

Asynchronous replication automatically adapts to the conditions by switching volume groups between journaling and cycling mode. If a volume group experiences too many journal-purges in a short period of time, it starts replicating by using cycling mode. This usually happens if the replication throughput is not high enough to sustain the write throughput for this volume group. Unlike Remote Copy Global Mirror with Change Volumes, a defined cycle period does not exist. Rather, replication will cycle frequently enough to ensure that the RPO that is defined by the replication policy is achieved (in the absence of errors). The system aims to ensure that all volume groups meet their defined RPO.

Conversely, if a volume group is found to be cycling too frequently it starts to replicate by using journaling mode. Journaling mode can be more resource-efficient, so it is the preferable option if the conditions can support it.

The transitions between the modes are transparent and are managed by the system. The mode that a volume group is currently using cannot be easily identified; the important aspect is the current recovery point, which is visible against the volume group. It is this metric that should be monitored. If the volume group exceeds the RPO that is defined on the replication policy, an alert is raised.

3.2.8 Status and recovery point reporting

The status of replication for a volume group is visible in the GUI on the Policies tab of the volume group view, or on the CLI using the **`lsvolumegroupreplication`** command.

The status includes information such as:

- Which system is currently production and which system is recovery, and therefore which direction replication is operating.
- Whether the systems are connected or disconnected.
- Whether replication is running or suspended.

- Information about the time of the recovery point that exists on the recovery copy, expressed either as a time behind production (if replication is running) or a fixed timestamp (if replication is not running).
- Whether the recovery point available for the volume group is within the recovery point objective (RPO) that is defined in the replication policy.

The recovery point can be tracked by using the statistics that are produced by the production system. It is available on the GUI, CLI, or REST API on both systems. This value is updated periodically and is rounded up to the nearest second. More granular reporting is only available within the XML statistics that can be retrieved by external monitoring tools.

> **Note:** A running recovery point of zero does not guarantee that the copies are identical as the value is updated periodically.

For more information on XML statistics, see the IBM FlashSystem® 9x00 documentation at: https://www.ibm.com/docs/en/flashsystem-9x00.

If a volume group exceeds the RPO of the associated replication policy, then an alert is generated in the event log. Depending on your notification settings, this can trigger a notification to be sent by email, **`syslog`**, or SNMP alerting.

If replication is suspended due to an error, the location status can be used to identify which system is unhealthy and requires attention. Replication automatically restarts when both systems are healthy.

3.3 Replication with snapshots, FlashCopy, and Safeguarded Copy

IBM Storage Virtualize version 8.5.1 introduced Volume Group Snapshots that operate on the same volume group definition as policy-based replication. With this feature, taking a snapshot of a volume group is a single GUI or CLI operation and can be performed on production, recovery, or independent volume groups. If a snapshot is taken of the recovery volume group, then the image of the snapshot is that of the current recovery point. Snapshot schedules can be configured to periodically take a snapshot of a volume group.

Safeguarded snapshots, which are used for protection against cyber-attacks, can be taken in the same way. It is also supported to use traditional FlashCopy maps to take a snapshot, clone, or backup of replicated volumes. However, traditional FlashCopy and Volume Group Snapshots cannot be mixed for the same volume.

As of version 8.5.2 of IBM Storage Virtualize, it is not possible to restore to a production volume using policy-based replication using either traditional FlashCopy or Volume Group Snapshots. Support might be added in a future version. It is also not possible to restore a recovery volume or volume group using traditional FlashCopy or Volume Group Snapshots as this would diverge the copy from the production system. The most current set of restrictions is documented in the IBM FlashSystem 9x00 documentation at: https://www.ibm.com/docs/en/flashsystem-9x00.

4

Implementing asynchronous policy-based replication

This section describes how to perform common tasks using the GUI or command-line interface (CLI).

4.1 Using the GUI

All aspects of replication can be managed using the GUI.

Upon creating the first partnership for policy-based replication, the management GUI displays a checklist to guide you through the remaining steps. The checklist includes the following items:

1. Completing the partnership setup by creating the partnership from the remote system.
2. Linking pools between systems, optionally using provisioning policies on each pool.
3. Creating a replication policy.
4. Creating a volume group and assigning a replication policy to the group.
5. Creating new volumes, or adding existing volumes, to the group.

The monitoring and management of replication is performed from the Volume Groups page.

4.1.1 Configuring policy-based replication using the setup checklist

To run the setup checklist, start the process of creating a 2-site partnership using the GUI. The example below is for a Fibre Channel (FC) partnership between two systems named 'london' and 'manchester'. The first step starts on the 'london' system, which is used as the production system for the first volume group. The 'manchester' system is used for the disaster recovery copy.

Creating a partnership

To create the partnership between systems, complete these steps:

1. On the first system, select **Copy Services** → **Partnerships and Remote Copy**.
2. Select **Create Partnership** and select **2-site partnership**.
3. Click **Continue**.
4. Select the remote system.
5. Ensure that the **Use Policy-Based Replication** checkbox is selected.
6. Enter the value, in megabits per second (Mbps), for the total bandwidth available between the two systems that can be used for replication. If all the bandwidth is available for policy-based replication, ensure the background copy rate (%) is set to **100**.
7. Click **Create**.

The checklist for configuring policy-based replication appears for the partnership (Figure 4-1 on page 27).

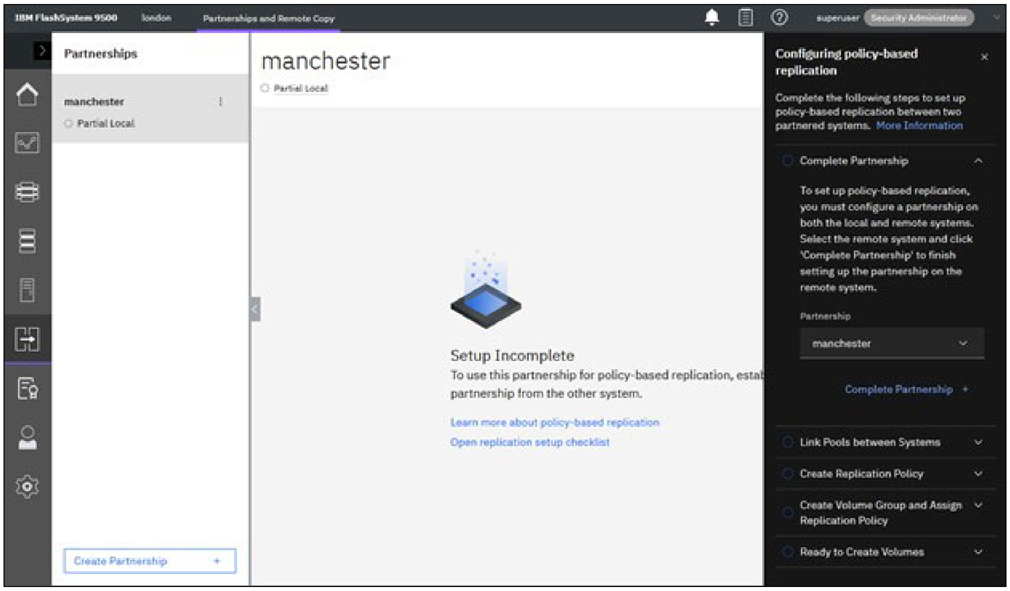
IBM FlashSystem 9500
london
Partnerships and Remote Copy
superuser
Security Administrator
Partnerships
manchester
Partial Local
Create Partnership +
manchester
Partial Local
Learn more about policy-based replication
Open replication setup checklist
Configuring policy-based replication
Complete Partnership
Partnership
manchester
Complete Partnership +
Link Pools between Systems
Create Replication Policy
Create Volume Group and Assign Replication Policy
Ready to Create Volumes

Setup Incomplete

To use this partnership for policy-based replication, estab partnership from the other system.

Complete the following steps to set up policy-based replication between two partnered systems. More Information

To set up policy-based replication, you must configure a partnership on both the local and remote systems. Select the remote system and click 'Complete Partnership' to finish setting up the partnership on the remote system.

Figure 4-1 Checklist for configuring policy-based replication

Select **Complete Partnership** in the checklist to open a new browser tab for the GUI of the remote system. After you sign in, the GUI launches to the Partnerships and Remote Copy page.

Follow the same steps that were used on the first system to create a partnership between the systems.

After the partnership is created from the second system, its status changes to **Configured** and the partnership indicates that it is ready for use with policy-based replication.

The checklist for the partnership appears again, which indicates that the partnership step is complete (Figure 4-2).

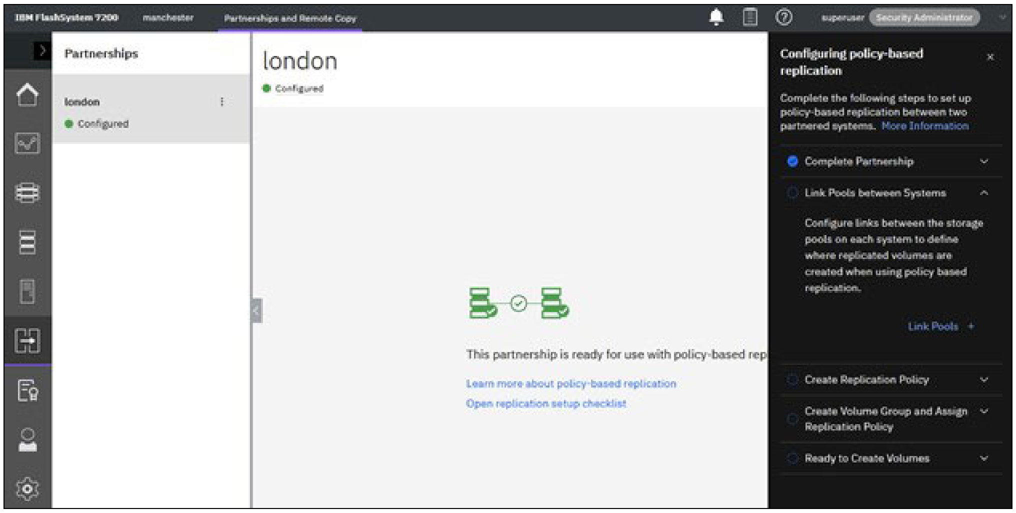
IBM FlashSystem 7200
manchester
Partnerships and Remote Copy
superuser
Security Administrator
Partnerships
london
Configured
london
Configured
Learn more about policy-based replication
Open replication setup checklist
Configuring policy-based replication
Complete Partnership
Link Pools between Systems
Link Pools +
Create Replication Policy
Create Volume Group and Assign Replication Policy
Ready to Create Volumes

This partnership is ready for use with policy-based rep

Complete the following steps to set up policy-based replication between two partnered systems. More Information

Configure links between the storage pools on each system to define where replicated volumes are created when using policy based replication.

Figure 4-2 Checklist for the partnership

At this point, you can either continue to use the GUI on this system to configure pool links between the systems or you can close this tab and return to the GUI of the first system.

Linking pools between systems

If the storage pools already exist on the production and recovery systems, you can add a link between the pools from either system. Select **Link Pools** in the setup checklist to add links between the pools.

With policy-based replication, storage pool links define which pool on the recovery system stores the recovery copy of the volume, based on the pool where the production volume exists.

In this example workflow, additional pools are created in the first system. Provisional policies are assigned to the pools before they are linked.

Return to the first system and select **Pools** → **Pools** to view the existing storage pools. Create new pools and add storage capacity to the pools as necessary.

Optionally, child pools can be created for the same parent pool array. Different provisioning policies for different capacity savings can be used for each child pool. When a child pool is created on the first system, select the **Child pool** option. The **Linked child pool** option can be used later in the setup process, when the corresponding linked pools are created on the second system.

A provisioning policy can be assigned to a pool to simplify volume creation. When a policy is assigned to a pool, any volumes that are created in the pool use the capacity savings methods that are defined in the policy. Provisioning policies can be managed from the **Policies** → **Provisioning Policies** entry in the navigation menu.

To assign a provisioning policy to an existing pool, complete these steps in the Pools view:

1. Select the pool.
2. Select **Actions** → **Assign Provisioning Policy**.
3. Select the provisioning policy to be assigned to the pool.
4. Click **Assign**.

In the example in Figure 4-3, two child pools are created with different provisioning policies.

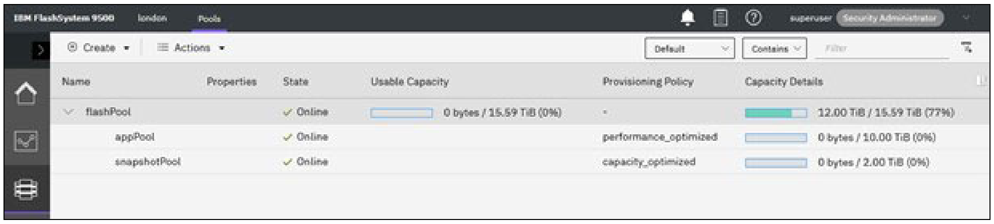

Name	Properties	State	Usable Capacity	Provisioning Policy	Capacity Details
flashPool		✓ Online	0 bytes / 15.59 TiB (0%)	-	12.00 TiB / 15.59 TiB (77%)
appPool		✓ Online		performance_optimized	0 bytes / 10.00 TiB (0%)
snapshotPool		✓ Online		capacity_optimized	0 bytes / 2.00 TiB (0%)

Figure 4-3 Two child pools created with different provisioning policies

In this example workflow, we assume that only the parent pool is created on the second system and the child pools still need to be created there. The GUI on the second system can be used to create and link pools to the first system. The management GUI simplifies the process of creating and linking a child pool on the recovery system in a single step. This allows the new pool to be created with attributes that match the remote pool without needing to log into the other system.

To create a linked child pool, complete these steps:

1. Select the parent pool.
2. Select **Actions** → **Create Child Pool**.

3. Select **Linked child pool**. The remote system information is on the left and the local system information is on the right.
4. On the left, select the remote system to which the new pool should be linked.
5. On the left, select the remote pool to which the new pool should be linked.
6. On the right, confirm that the desired local parent pool is selected.
7. Enter a name for the new pool.
8. If necessary, specify the capacity of the new child pool and whether the pool should use encryption.
9. On the right, select the provisioning policy for the new pool.
10. Click **Create** (Figure 4-4).

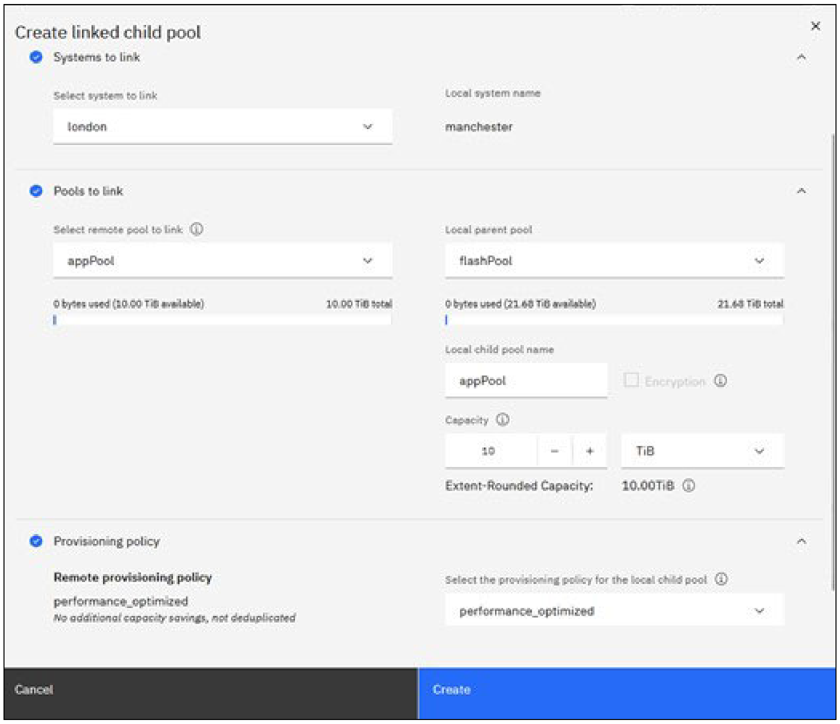

Create linked child pool

Systems to link

Select system to link: london

Local system name: manchester

Pools to link

Select remote pool to link: appPool

0 bytes used (10.00 TiB available) 10.00 TiB total

Local parent pool: flashPool

0 bytes used (21.68 TiB available) 21.68 TiB total

Local child pool name: appPool

□ Encryption

Capacity: 10 TiB

Extent-Rounded Capacity: 10.00TiB

Provisioning policy

Remote provisioning policy

performance_optimized

No additional capacity savings, not deduplicated

Select the provisioning policy for the local child pool: performance_optimized

Figure 4-4 Create linked child pool

Linked pools are indicated by a link icon in the **Properties** column on the Pools page (Figure 4-5).

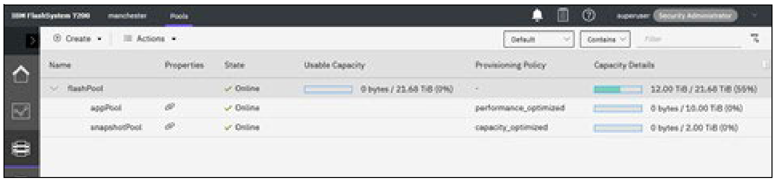

Name	Properties	State	Usable Capacity	Provisioning Policy	Capacity Details
flashPool		✓ Online	0 bytes / 21.68 TiB (0%)	-	12.00 TiB / 21.68 TiB (55%)
appPool	🔗	✓ Online		performance_optimized	0 bytes / 10.00 TiB (0%)
snapshotPool	🔗	✓ Online		capacity_optimized	0 bytes / 2.00 TiB (0%)

Figure 4-5 Linked pools

To display the details of the pool links that are configured for a pool, complete these steps:

1. Select the pool.
2. Select **Actions** → **Properties**.
3. Review the pool link information.

4. Click **Close** (Figure 4-6).

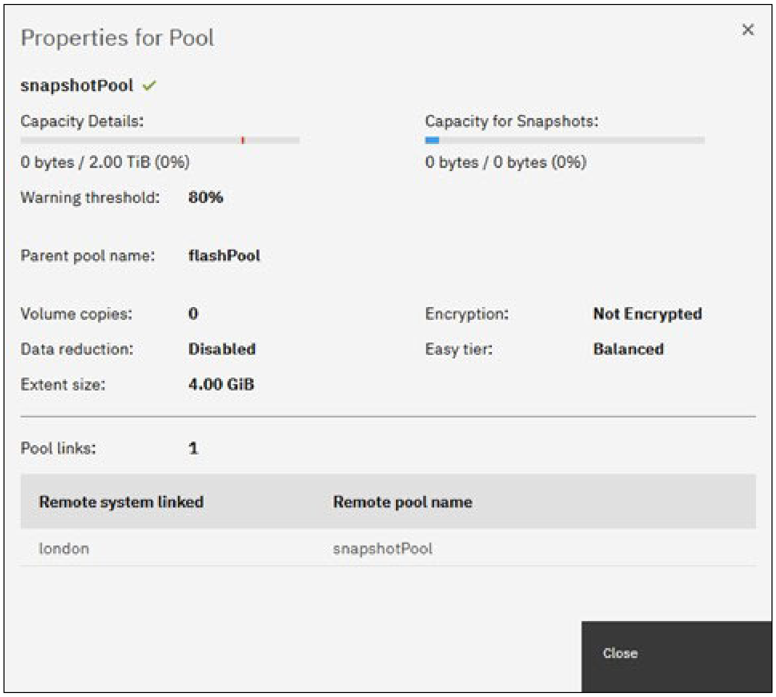

Properties for Pool

snapshotPool ✓

Capacity Details: 0 bytes / 2.00 TiB (0%)

Capacity for Snapshots: 0 bytes / 0 bytes (0%)

Warning threshold: 80%

Parent pool name: flashPool

Volume copies: 0

Encryption: Not Encrypted

Data reduction: Disabled

Easy tier: Balanced

Extent size: 4.00 GiB

Pool links: 1

Remote system linked	Remote pool name
london	snapshotPool

Figure 4-6 Properties for pool

After pools are linked between systems, the replication setup checklist updates to show that the step is complete. If the checklist is not visible, select **Open replication setup checklist** on the partnership view.

You do not need to repeat the process on the remote system; the pool links are updated automatically on the partnered systems.

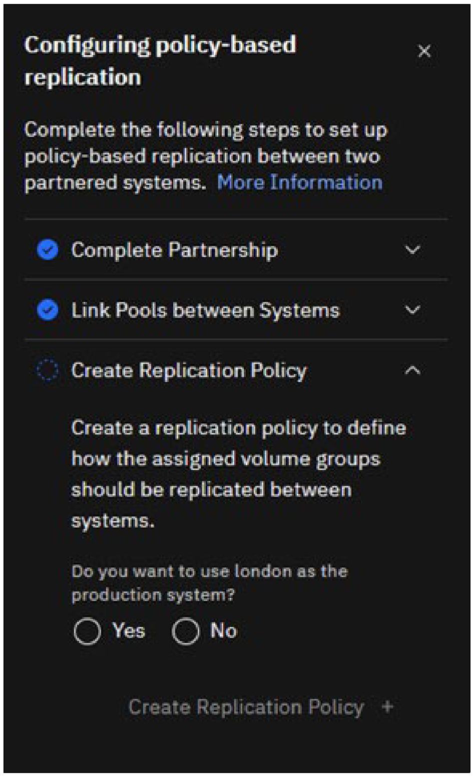

Configuring policy-based replication ×

Complete the following steps to set up policy-based replication between two partnered systems. More Information

- Complete Partnership
- Link Pools between Systems
- Create Replication Policy

Create a replication policy to define how the assigned volume groups should be replicated between systems.

Do you want to use london as the production system?

○ Yes ○ No

Create Replication Policy +

Figure 4-7 Configuring policy-based replication

Creating a replication policy

To create a replication policy, open the replication setup checklist on the system that is to be used as the production system for the first volume group. In this example workflow, the production system is the 'london' system.

1. Select **Yes** to use the local system as the production system.
2. Click **Create Replication Policy**.

The **Replication Policies** page is displayed and the **Create Replication Policy** task is launched.

To create a replication policy, complete these steps:

1. Enter a name for the replication policy. This name cannot be changed after the policy is created and must be unique on both systems.
2. Select the **Topology** to represent how the systems are organized and the type of replication between systems. In this example, **2 Site**, **Asynchronous** is preselected.
3. On the left under **Location 1**, select the first system to be used in the policy. This system appears on the left when managing and monitoring replication.
4. On the right under **Location 2**, select the second system to be used in the policy. This system appears on the right when managing and monitoring replication.
5. If either system has multiple I/O groups, select the I/O group that will be the caching I/O group for all volumes that are replicated with this policy.
6. Under **Recovery point objective (RPO)**, enter the recovery point objective for the policy, in minutes. An alert is sent if the recovery point exceeds this value.
7. Click **Create** (Figure 4-8 on page 32).

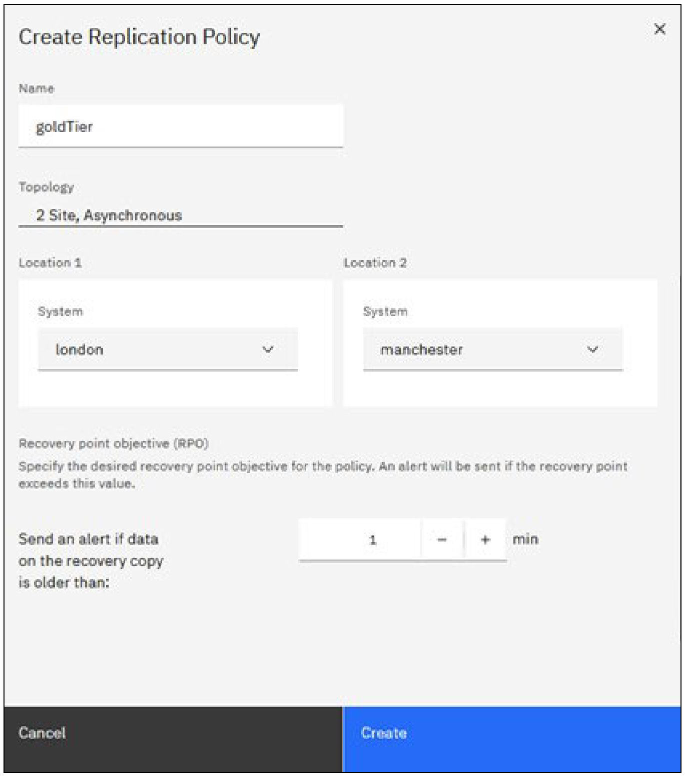

Create Replication Policy

Name: goldTier

Topology: 2 Site, Asynchronous

Location 1 — System: london

Location 2 — System: manchester

Recovery point objective (RPO)

Specify the desired recovery point objective for the policy. An alert will be sent if the recovery point exceeds this value.

Send an alert if data on the recovery copy is older than: 1 min

Figure 4-8 Create Replication Policy

The new replication policy is displayed and the replication setup checklist is updated to show that the step is complete. The system automatically creates the same replication policy on the remote system.

If required, additional replication policies can be created on this page (Figure 4-9) before you continue with the next step in "Creating a volume group and assigning a replication policy" on page 33.

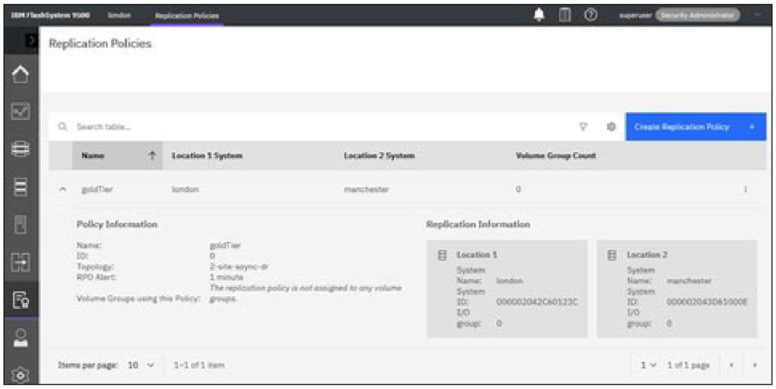

Replication Policies

Name	Location 1 System	Location 2 System	Volume Group Count
goldTier	london	manchester	0

Policy Information

Name: goldTier
ID: 0
Topology: 2-site-async-dr
RPO Alert: 1 minute
Volume Groups using this Policy: *The replication policy is not assigned to any volume groups.*

Replication Information

Location 1
System Name: london
System ID: 000002042C60123C
I/O group: 0

Location 2
System Name: manchester
System ID: 000002043D61000E
I/O group: 0

Items per page: 10 1–1 of 1 item 1 1 of 1 page

Figure 4-9 Replication Policies

If the replication setup checklist is not visible, it can be launched by clicking the '**?**' icon in the top menu bar and selecting **Replication Setup** (Figure 4-10 on page 33).

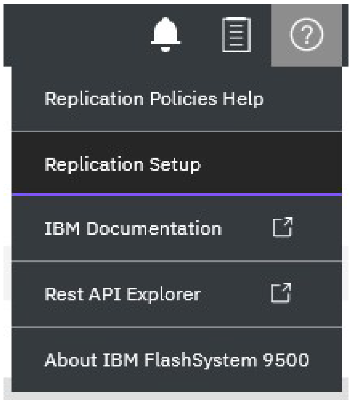

Figure 4-10 Replication Setup

The Configuring policy-based replication page is displayed (Figure 4-11).

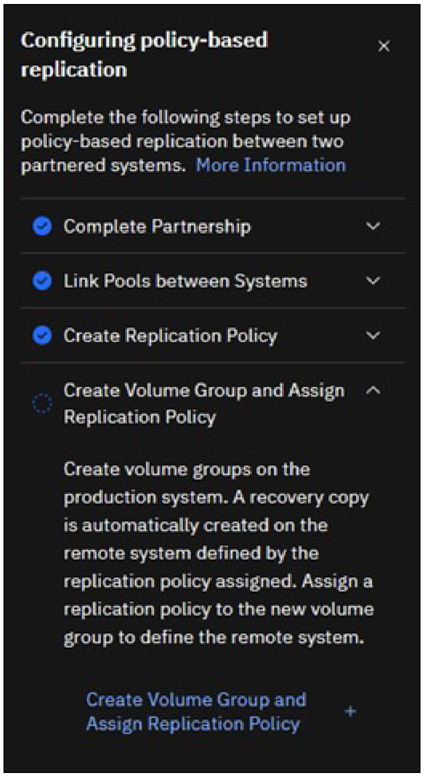

Configuring policy-based replication

Complete the following steps to set up policy-based replication between two partnered systems. More Information

Complete Partnership
Link Pools between Systems
Create Replication Policy
Create Volume Group and Assign Replication Policy

Create volume groups on the production system. A recovery copy is automatically created on the remote system defined by the replication policy assigned. Assign a replication policy to the new volume group to define the remote system.

Figure 4-11 Configuring policy-based replication

Creating a volume group and assigning a replication policy

To create a volume group and assign a replication policy on the system that is used as the production system, complete these steps:

1. Select **Create Volume Group and Assign Replication Policy** in the checklist.
2. In the **Create Volume Group** task, enter a name for the volume group. After a replication policy is assigned to the group, this name cannot be changed and must be unique on both systems.
3. Click **Create Empty Group**.

4. In the **Assign Replication Policy** task, select the replication policy to assign to the volume group. A preview that indicates the direction of replication is displayed.
5. Click **Assign** (Figure 4-12).

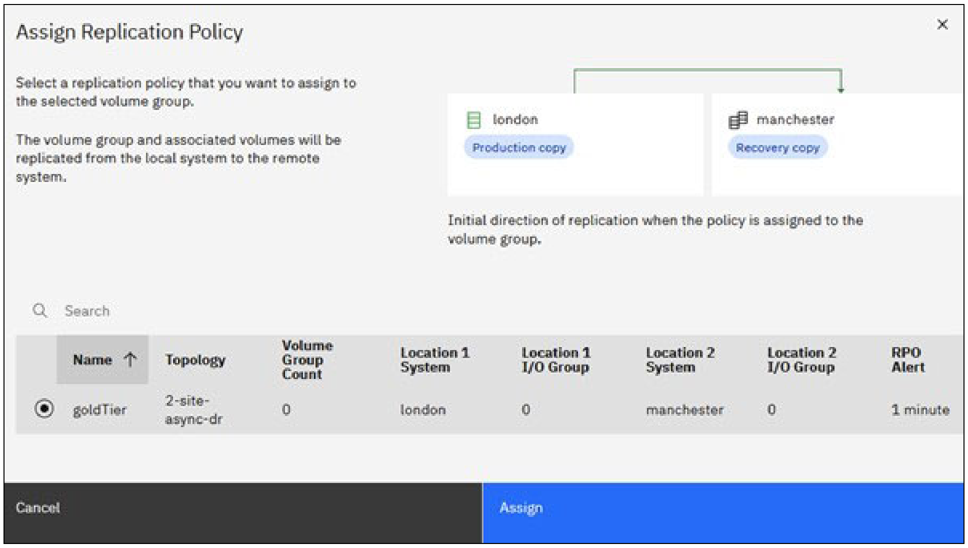

Assign Replication Policy

Select a replication policy that you want to assign to the selected volume group.

The volume group and associated volumes will be replicated from the local system to the remote system.

Initial direction of replication when the policy is assigned to the volume group.

	Name ↑	Topology	Volume Group Count	Location 1 System	Location 1 I/O Group	Location 2 System	Location 2 I/O Group	RPO Alert
●	goldTier	2-site-async-dr	0	london	0	manchester	0	1 minute

Figure 4-12 Assign Replication Policy

The new volume group is created and the replication policy is assigned (Figure 4-13).

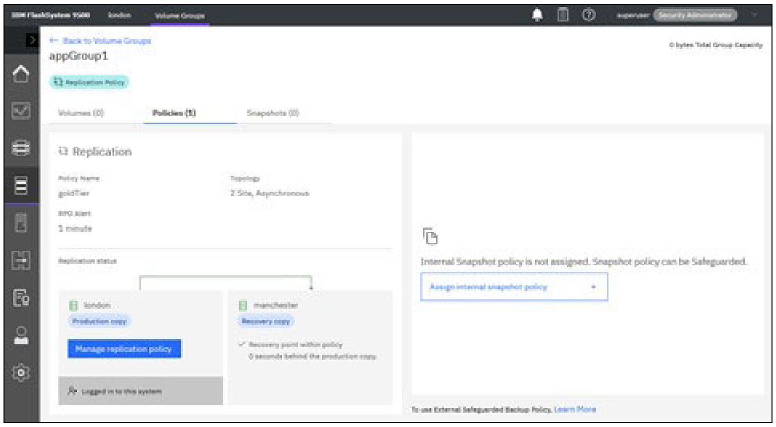

appGroup1

Policy Name: goldTier

Topology: 2 Site, Asynchronous

RPO Alert: 1 minute

Replication status

✓ Recovery point within policy

0 seconds behind the production copy.

Internal Snapshot policy is not assigned. Snapshot policy can be Safeguarded.

To use External Safeguarded Backup Policy, Learn More

Figure 4-13 New volume group is created and the replication policy is assigned

The replication status indicates that replication is running from 'london' to 'manchester' for this volume group.

The replication setup checklist is updated to show the step is complete and policy-based replication is ready to use (Figure 4-14 on page 35).

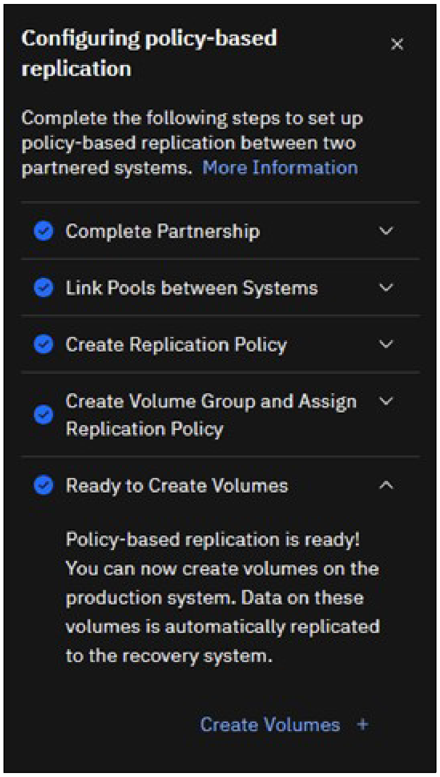

Configuring policy-based replication

Complete the following steps to set up policy-based replication between two partnered systems. More Information

- Complete Partnership
- Link Pools between Systems
- Create Replication Policy
- Create Volume Group and Assign Replication Policy
- Ready to Create Volumes

Policy-based replication is ready! You can now create volumes on the production system. Data on these volumes is automatically replicated to the recovery system.

Create Volumes +

Figure 4-14 Replication setup checklist is updated

The **Create Volumes** link in the checklist opens the Volumes page, where new volumes can be created in the volume group or existing volumes can be moved into the group.

4.1.2 Creating new volumes in a volume group

Creating new volumes directly in a volume group allows the initial synchronization to be optimized, even if the systems are disconnected.

To create new volumes in a volume group, complete these steps on the Volumes page:

1. Click **Create Volumes**.
2. Select the volume group in which to create the volume.
3. Continue to define the volume properties and create the volumes.

Alternatively, complete these steps on the Volume Groups page:

1. Select the volume group in which to create the volumes.
2. Select the **Volumes** tab for that volume group.
3. Select **Actions** → **Create New Volumes**.
4. Continue to define the volume properties and create the volumes.

When new volumes are created in the production copy of a volume group, the system automatically creates the recovery copies of those volumes and configures replication according to the assigned replication policy.

4.1.3 Moving volumes into a volume group

To add existing volumes to a volume group, complete these steps on the Volumes page:

1. Select the volumes to be added to the volume group.
2. Click **Actions** → **Add to Volume Group**.
3. Select the volume group to add the volumes to.
4. Click **Add volumes**.

Alternatively, complete these steps on the Volume Groups page:

1. Select the volume group to add the volumes to.
2. Select the **Volumes** tab for the volume group.
3. Select **Actions** → **Add Existing Volumes**.
4. Select the volumes to be added to the volume group.
5. Click A**dd Existing Volumes**.

When volumes are added to the production copy of a volume group, the system automatically creates copies of those volumes in the recovery copy of the volume group. The system configures replication according to the policy that is assigned to the group.

4.1.4 Viewing volumes in a volume group

To view the volumes in a volume group, complete these steps on the Volumes page:

1. Select **Volumes** > **Volumes Groups**.
2. Select the volume group to view.
3. Select the **Volumes** tab for the volume group.

Figure 4-15 shows an example view on the system with the production copy of a volume group.

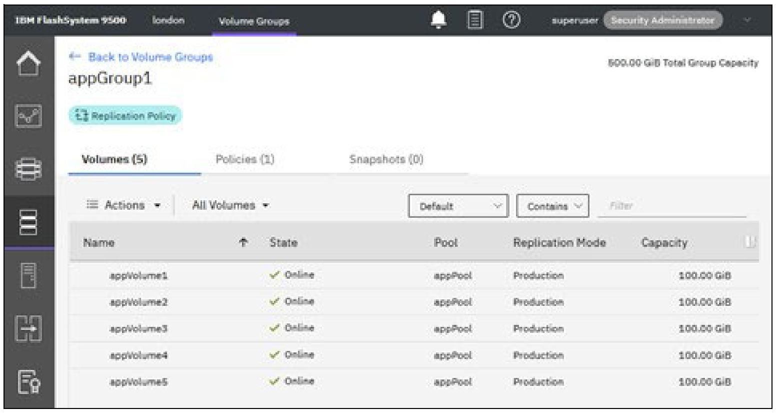

Name	State	Pool	Replication Mode	Capacity
appVolume1	✓ Online	appPool	Production	100.00 GiB
appVolume2	✓ Online	appPool	Production	100.00 GiB
appVolume3	✓ Online	appPool	Production	100.00 GiB
appVolume4	✓ Online	appPool	Production	100.00 GiB
appVolume5	✓ Online	appPool	Production	100.00 GiB

Figure 4-15 Example view on the system with the production copy of a volume group

Figure 4-16 on page 37 shows an example view on the system with the recovery copy of a volume group.

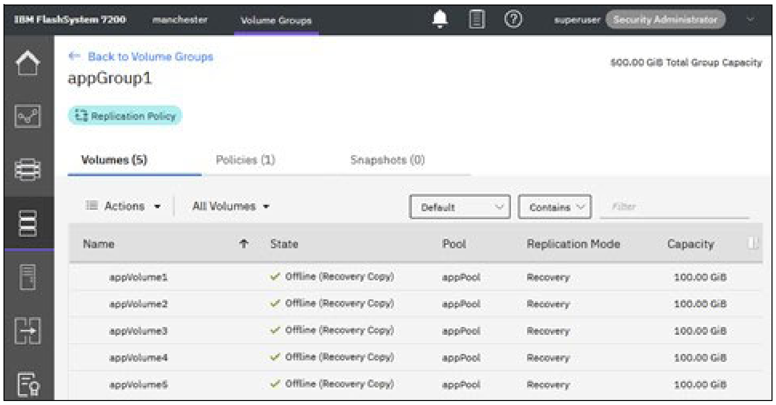

appGroup1

500.00 GiB Total Group Capacity

Name	State	Pool	Replication Mode	Capacity
appVolume1	✓ Offline (Recovery Copy)	appPool	Recovery	100.00 GiB
appVolume2	✓ Offline (Recovery Copy)	appPool	Recovery	100.00 GiB
appVolume3	✓ Offline (Recovery Copy)	appPool	Recovery	100.00 GiB
appVolume4	✓ Offline (Recovery Copy)	appPool	Recovery	100.00 GiB
appVolume5	✓ Offline (Recovery Copy)	appPool	Recovery	100.00 GiB

Figure 4-16   Example view on the system with the recovery copy of a volume group

4.1.5 Deleting volumes in a volume group

To delete a volume in a volume group that uses policy-based replication, delete the volume on the production system. The volume is automatically deleted on the recovery system when the recovery point updates to one that does not include the volume.

4.1.6 Moving volumes out of a volume group

To move a volume out of a volume group, complete these steps on the Volumes page:

1. Select **Volumes** → **Volume Groups**.
2. Select the volume group that contains the volume to be moved.
3. Select the **Volumes** tab for the volume group and select the volume to be moved.
4. Select **Actions** → **Remove from Volume Group**.
5. Click **Remove**.

4.1.7 Managing provisioning policies

Provisioning policies can be created, renamed, and deleted on the Provisioning Policies page (Figure 4-17).

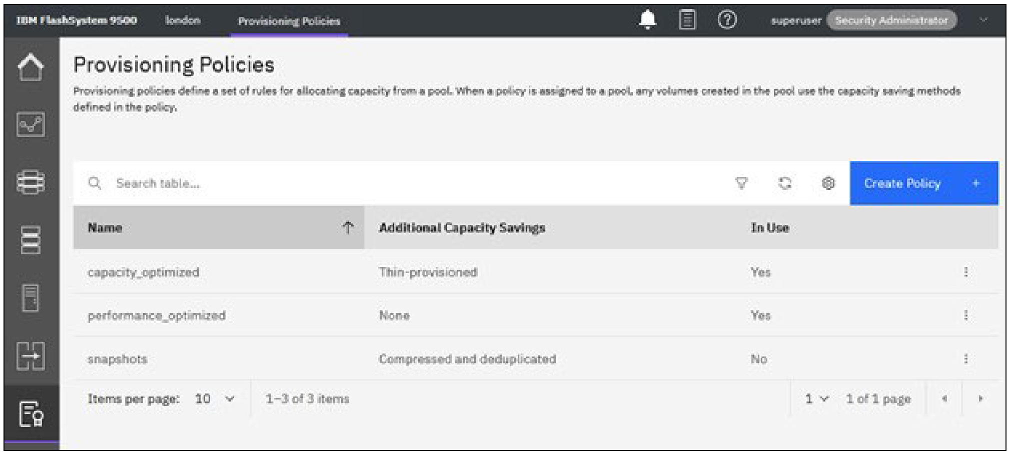

Provisioning Policies

Provisioning policies define a set of rules for allocating capacity from a pool. When a policy is assigned to a pool, any volumes created in the pool use the capacity saving methods defined in the policy.

Name	Additional Capacity Savings	In Use
capacity_optimized	Thin-provisioned	Yes
performance_optimized	None	Yes
snapshots	Compressed and deduplicated	No

Figure 4-17   Provisioning Policies

Provisioning policies for storage pools can be managed on the Pools page **(**Figure 4-18).

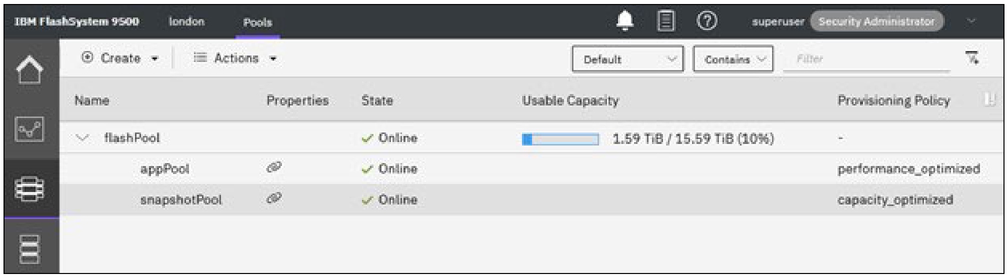

Name	Properties	State	Usable Capacity	Provisioning Policy
flashPool		✓ Online	1.59 TiB / 15.59 TiB (10%)	-
appPool		✓ Online		performance_optimized
snapshotPool		✓ Online		capacity_optimized

Figure 4-18 Provisioning policies for storage pools

To assign a provisioning policy to a pool, complete these steps on the Provisioning Policies page:

1. Select the pool.
2. Click **Actions** → **Assign Provisioning Policy**.
3. Select the provisioning policy to assign to the pool.
4. Click **Assign**.

To unassign a provisioning policy from a pool or to assign an alternative policy, complete these steps on the Provisioning Policies page:

1. Select the pool.
2. Click **Actions** → **Manage Provisioning Policy**.
3. Select the new provisioning policy to assign to the pool, or select **Remove policy**.
4. Click **Apply changes**.

Any new volumes that are created in the pool are automatically created by using the capacity savings methods that are defined in the provisioning policy. Volumes that are already existing in the pool are unchanged.

4.1.8 Managing pool links

Storage pool links define which pool on the recovery system stores the recovery copy of the volume, based on the pool where the production volume exists. The pools on the recovery system require links to the pools on the production system. Links are required on any pool that contains volumes that use policy-based replication. Policy-based replication requires at least one pool to be linked on each system.

Use the Pools page to create and manage storage pool links.

If neither of the pools to be linked are already linked to other systems, the link can be configured from either system. If a pool on one of the systems contains links to another partnered system, you must add the new link from the unlinked system.

To link pools for policy-based replication, complete these steps on the Pools page:

1. Select the pool.
2. Select **Actions** → **Add Pool Link for Replication**.

   The remote system information is shown on the left and the local system information is on the right.

3. On the left, select the remote system to which the new pool should be linked.
4. On the left, select the remote pool to which the new pool should be linked.
5. On the right, confirm that the desired local pool is selected.
6. Optionally, assign provisioning policies to each of the pools, or select the **Use standard provisioning for replicated volumes** checkbox.
7. Click **Apply changes** (Figure 4-19).

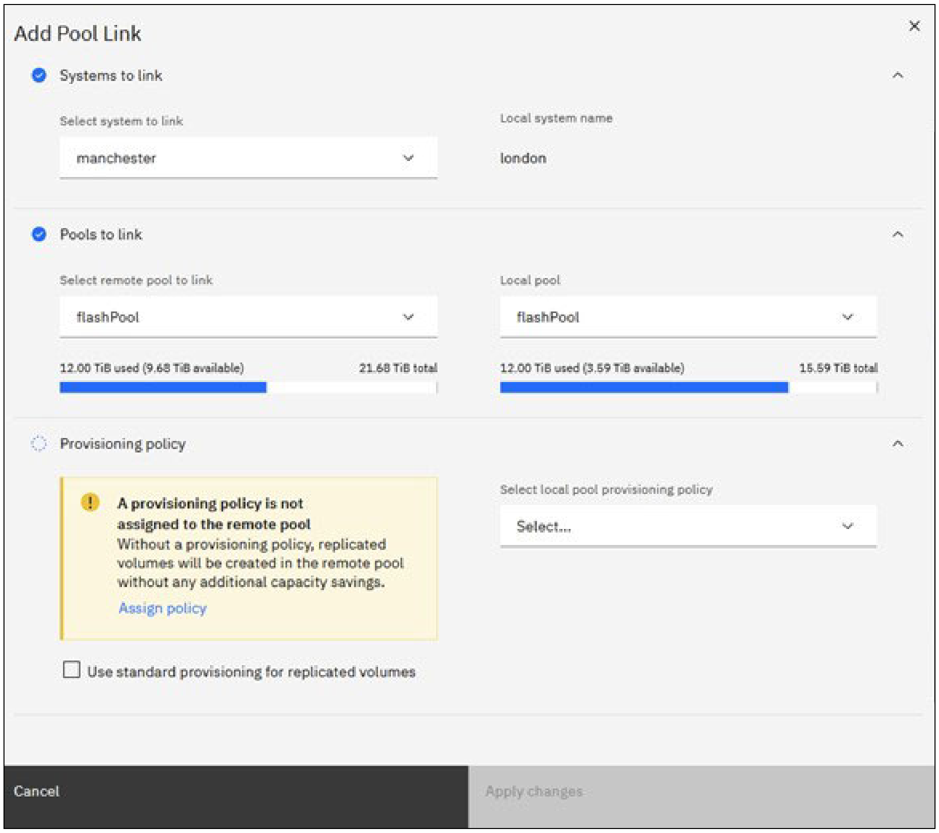

Select system to link: manchester

Local system name: london

Select remote pool to link: flashPool

12.00 TiB used (9.68 TiB available) 21.68 TiB total

Local pool: flashPool

12.00 TiB used (3.59 TiB available) 15.59 TiB total

A provisioning policy is not assigned to the remote pool
Without a provisioning policy, replicated volumes will be created in the remote pool without any additional capacity savings.
Assign policy

Select local pool provisioning policy: Select...

□ Use standard provisioning for replicated volumes

Figure 4-19 Add Pool Link

You do not need to repeat the process on the remote system; the pool links are updated automatically on the partnered systems.

Linked pools are indicated by a link icon in the pool **Properties** on the Pools page (Figure 4-20).

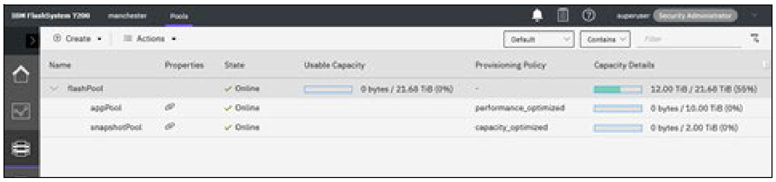

Name	Properties	State	Usable Capacity	Provisioning Policy	Capacity Details
flashPool		✓ Online	0 bytes / 21.68 TiB (0%)	-	12.00 TiB / 21.68 TiB (55%)
appPool	🔗	✓ Online		performance_optimized	0 bytes / 10.00 TiB (0%)
snapshotPool	🔗	✓ Online		capacity_optimized	0 bytes / 2.00 TiB (0%)

Figure 4-20 Pool Properties

To display the details of the pool links that are configured for a pool, complete these steps:

1. Select the pool.
2. Select **Actions** → **Properties**.

3. Review the pool link information, then click **Close** (Figure 4-21).

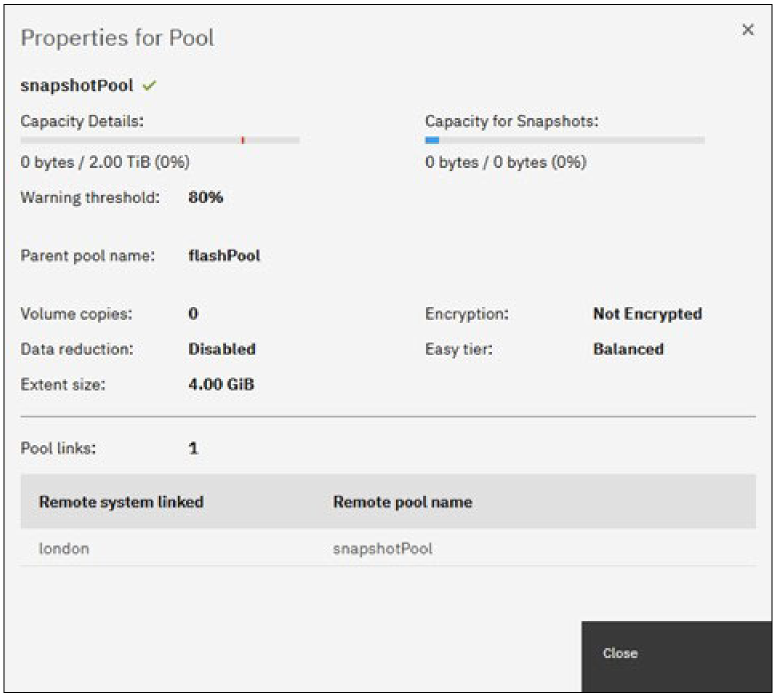

Properties for Pool

snapshotPool ✓

Capacity Details: 0 bytes / 2.00 TiB (0%)

Capacity for Snapshots: 0 bytes / 0 bytes (0%)

Warning threshold: 80%

Parent pool name: flashPool

Volume copies: 0

Encryption: Not Encrypted

Data reduction: Disabled

Easy tier: Balanced

Extent size: 4.00 GiB

Pool links: 1

Remote system linked	Remote pool name
london	snapshotPool

Figure 4-21 Properties for Pool

4.1.9 Managing replication policies

To create a replication policy, complete these steps:

1. Select **Policies** → **Replication Policies**.
2. On the Replication Policies page, select **Create Replication Policy**.
3. Enter a name for the replication policy. This name cannot be changed after the policy is and must be unique on both systems.
4. Select the **Topology** to represent how the systems are organized and the type of replication between systems. In this example, **2 Site, Asynchronous** is preselected.
5. On the left under **Location 1**, select the first system to be used in the policy. This system appears on the left when managing and monitoring replication.
6. On the right under **Location 2**, select the second system to be used in the policy. This system appears on the right when managing and monitoring replication.
7. If either system has multiple I/O groups, select the I/O group that is the caching I/O group for all volumes that are replicated with this policy.
8. Under **Recovery point objective (RPO)**, enter the recovery point objective for the policy, in minutes. An alert is sent if the recovery point exceeds this value.
9. Click **Create** (Figure 4-22 on page 41).

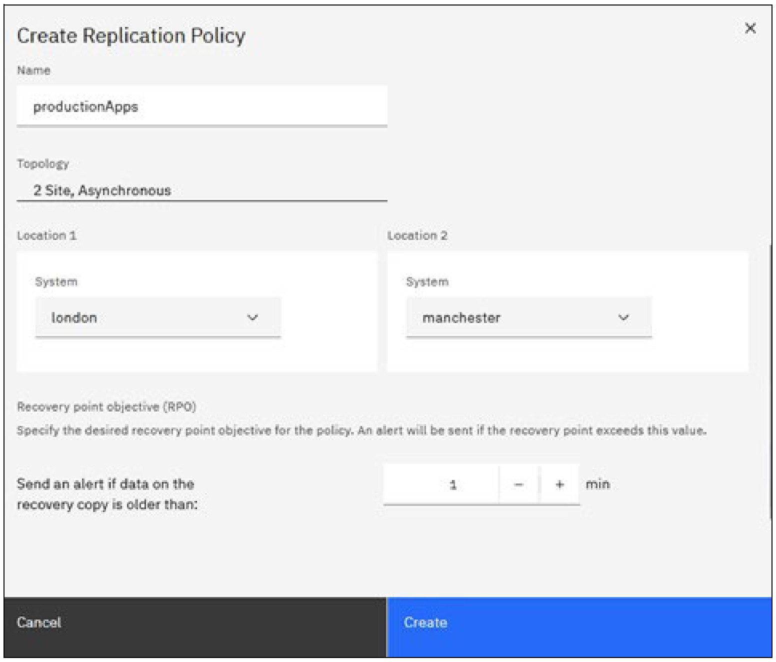

Create Replication Policy

Name

productionApps

Topology

2 Site, Asynchronous

Location 1	Location 2
System london	System manchester

Recovery point objective (RPO)

Specify the desired recovery point objective for the policy. An alert will be sent if the recovery point exceeds this value.

Send an alert if data on the recovery copy is older than: 1 min

Figure 4-22 Create Replication Policy

The replication policies are displayed. The system automatically creates the new replication policy on the remote system (Figure 4-23).

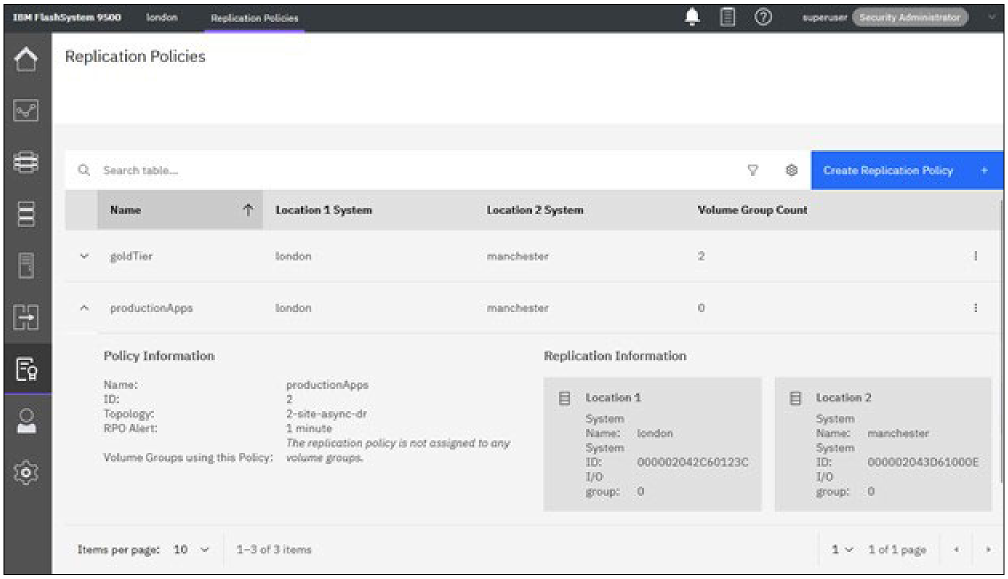

Replication Policies

Name	Location 1 System	Location 2 System	Volume Group Count
goldTier	london	manchester	2
productionApps	london	manchester	0

Policy Information

Name: productionApps
ID: 2
Topology: 2-site-async-dr
RPO Alert: 1 minute
Volume Groups using this Policy: *The replication policy is not assigned to any volume groups.*

Replication Information

Location 1	Location 2
System Name: london	System Name: manchester
System ID: 000002042C60123C	System ID: 000002043D61000E
I/O group: 0	I/O group: 0

Items per page: 10 1–3 of 3 items 1 1 of 1 page

Figure 4-23 Replication Policies

To delete a replication policy, complete these steps:

1. Select **Policies** → **Replication Policies**.
2. Select **Delete** from the overflow menu on the policy.

The system automatically removes the replication policy on the remote system when the systems are connected.

4.1.10 Viewing the replication status of a volume group

To view the replication status for the volume groups on the system:

Select **Volumes** → **Volumes Groups**.

The Volume Groups summary view shows whether replication is running and if the recovery point is within the objective that is specified in the replication policy (Figure 4-24).

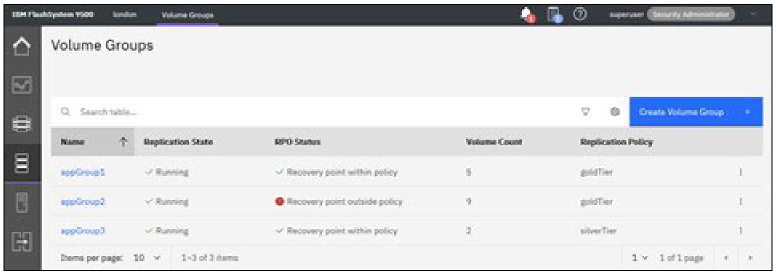

Volume Groups

Search table...

Name	Replication State	RPO Status	Volume Count	Replication Policy
appGroup1	✓ Running	✓ Recovery point within policy	5	goldTier
appGroup2	✓ Running	● Recovery point outside policy	9	goldTier
appGroup3	✓ Running	✓ Recovery point within policy	2	silverTier

Items per page: 10 1–3 of 3 items 1 1 of 1 page

Figure 4-24 Volume Groups summary view

To view detailed information for an individual volume group and manage the replication settings, click the volume group name. The replication status panel displays the replication mode of each system and information about the recovery point (Figure 4-25).

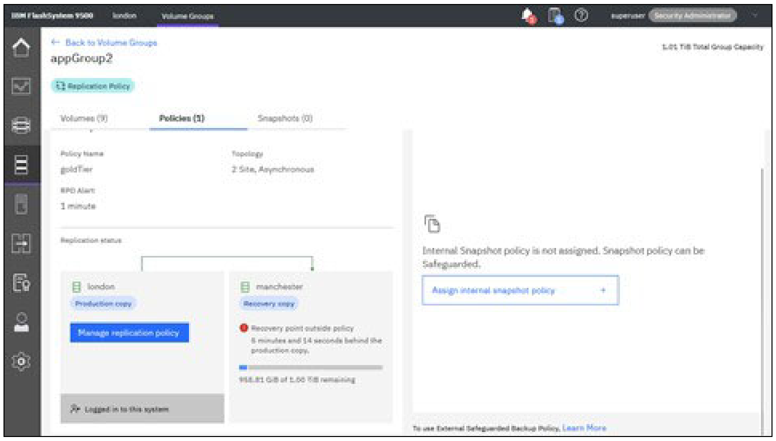

appGroup2

1.01 TiB Total Group Capacity

Policy Name: goldTier

Topology: 2 Site, Asynchronous

RPO Alert: 1 minute

Replication status

london — Production copy

manchester — Recovery copy

● Recovery point outside policy
6 minutes and 14 seconds behind the production copy.

958.81 GiB of 1.00 TiB remaining

Internal Snapshot policy is not assigned. Snapshot policy can be Safeguarded.

To use External Safeguarded Backup Policy, Learn More

Figure 4-25 Replication status panel

When the systems are connected, the same information is available on either system. When the systems are disconnected, information for only the local system is available.

Figure 4-26 on page 43 and Figure 4-27 on page 43 show example views on the system with the production copy of a volume group.

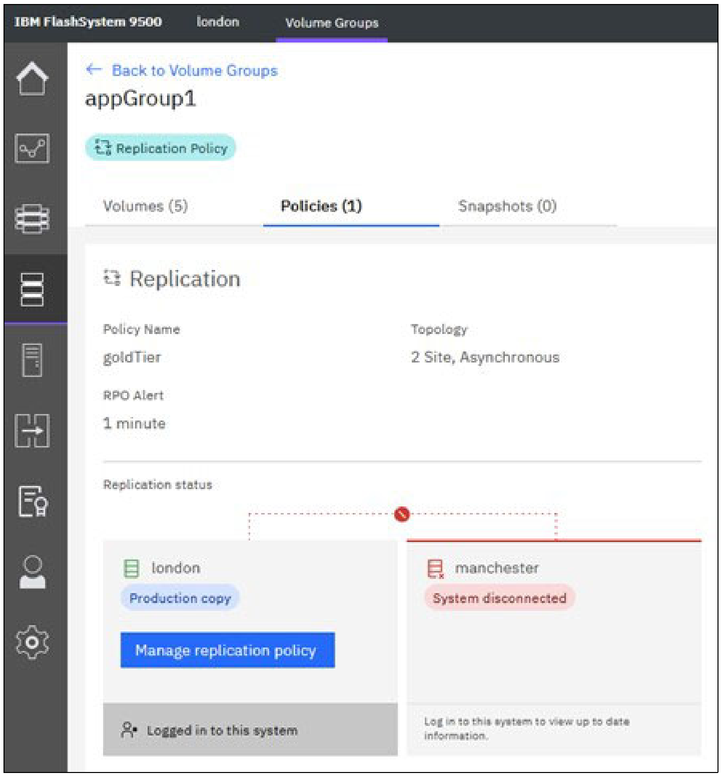

appGroup1

Policy Name: goldTier

Topology: 2 Site, Asynchronous

RPO Alert: 1 minute

Replication status

london: Production copy

Logged in to this system

manchester: System disconnected

Log in to this system to view up to date information.

Figure 4-26 Example view on the system with the production copy of a volume group -1

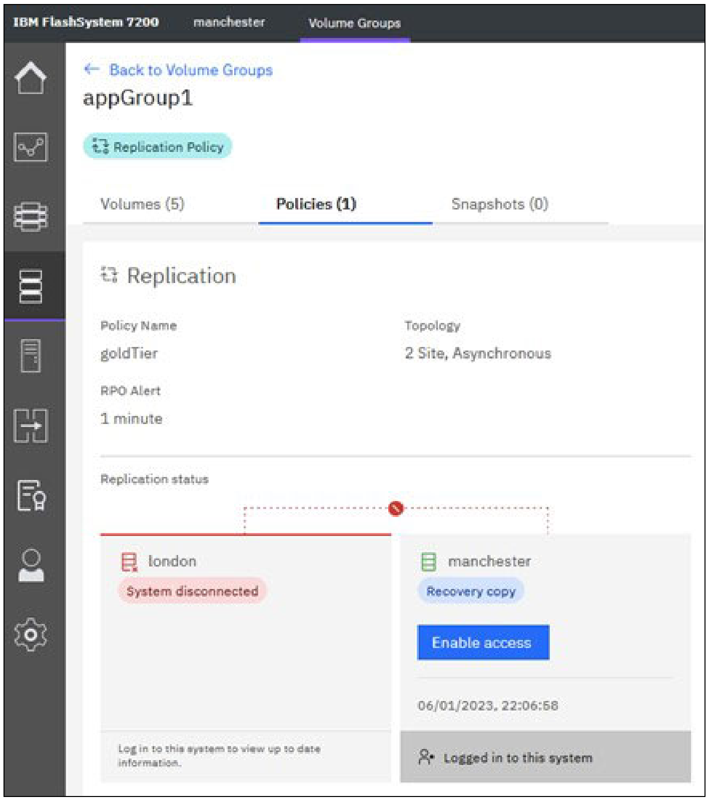

appGroup1

Policy Name: goldTier

Topology: 2 Site, Asynchronous

RPO Alert: 1 minute

Replication status

london: System disconnected

Log in to this system to view up to date information.

manchester: Recovery copy

06/01/2023, 22:06:58

Logged in to this system

Figure 4-27 Example view on the system with the production copy of a volume group -2

4.1.11 Assigning a replication policy to a volume group

Policy-based replication is configured on a volume group by assigning a replication policy to the group. The system where the policy is assigned is the production system, and the other system in the policy is the recovery system for the volume group.

To assign a replication policy to a volume group, complete these steps:

1. Select **Volumes** → **Volume Groups**.
2. Click the volume group name to view details for the group.
3. Click the **Policies** tab for the volume group.
4. Click **Assign replication policy**.
5. Select the replication policy to assign. A preview that indicates the direction of replication is displayed.
6. Click **Assign** (Figure 4-28).

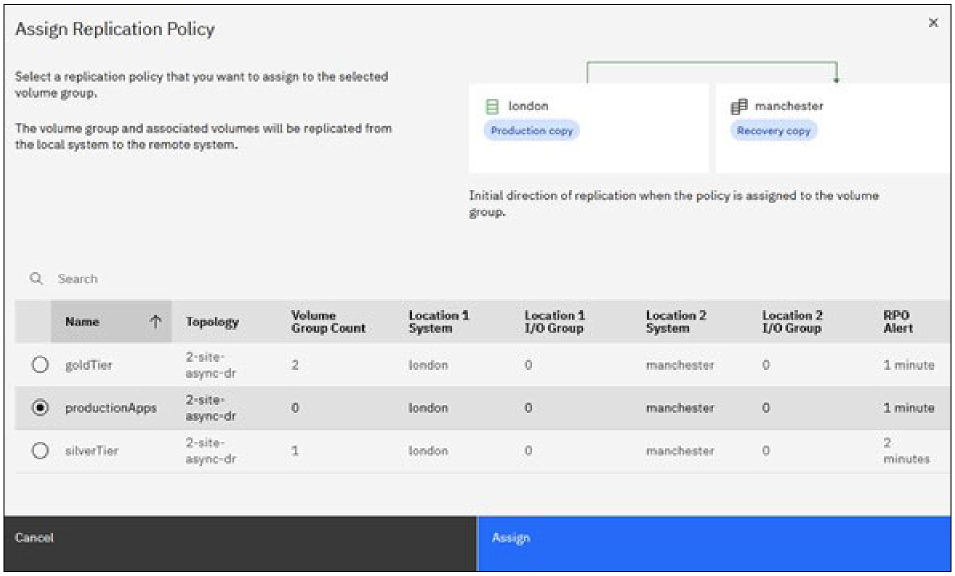

Assign Replication Policy

Select a replication policy that you want to assign to the selected volume group.

The volume group and associated volumes will be replicated from the local system to the remote system.

Initial direction of replication when the policy is assigned to the volume group.

	Name	Topology	Volume Group Count	Location 1 System	Location 1 I/O Group	Location 2 System	Location 2 I/O Group	RPO Alert
○	goldTier	2-site-async-dr	2	london	0	manchester	0	1 minute
●	productionApps	2-site-async-dr	0	london	0	manchester	0	1 minute
○	silverTier	2-site-async-dr	1	london	0	manchester	0	2 minutes

Figure 4-28 Assign Replication Policy

The system creates the recovery copy of the volume group and volumes on the remote system automatically.

If volumes are already in the volume group, the progress of the initial copy is displayed on the Replication status panel. When the initial copy completes, the recovery point is displayed (Figure 4-29 on page 45).

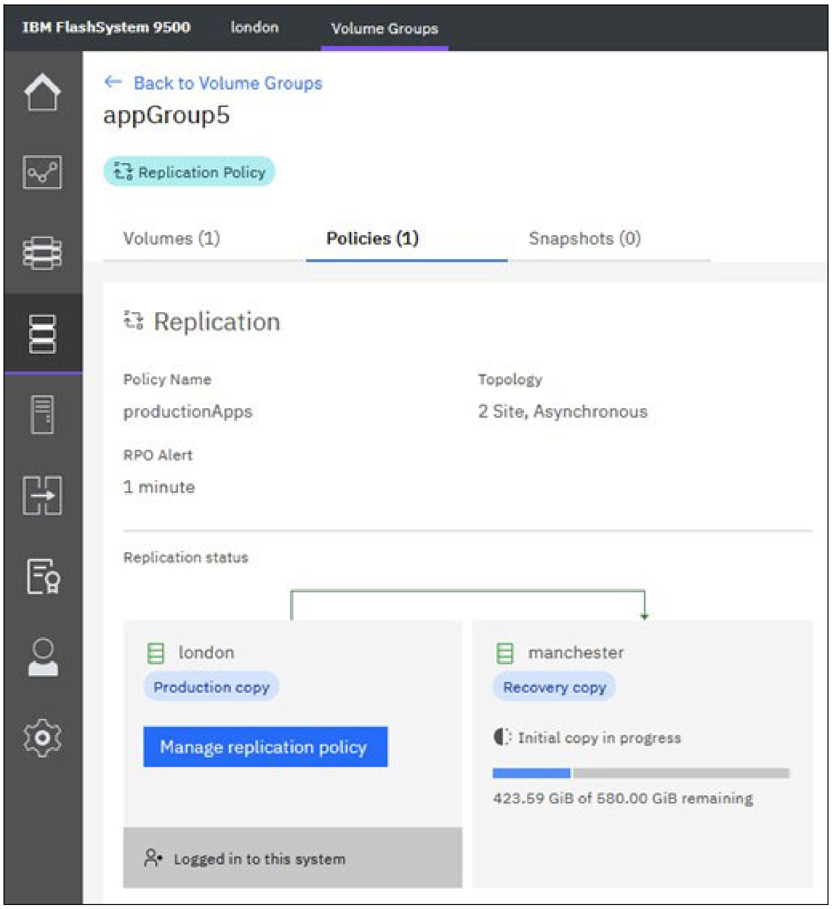

Policy Name
productionApps

Topology
2 Site, Asynchronous

RPO Alert
1 minute

Replication status

Initial copy in progress

423.59 GiB of 580.00 GiB remaining

Figure 4-29 Progress of the initial copy is displayed on the replication status panel

4.1.12 Replacing the replication policy assigned to a volume group

The attributes of a replication policy cannot be modified after the policy is created. However, a new policy that uses the same topology and locations can be created and assigned to the volume groups to replace the existing policy. When no volume groups use a replication policy, then that policy can be deleted.

To assign a new policy to a volume group, complete these steps using the GUI on the production system:

1. Select **Volumes** > **Volume Groups**.
2. Click on the volume group name to view details for the group.
3. Click the **Policies** tab for the volume group.
4. Click **Manage replication policy** (Figure 4-30 on page 46).

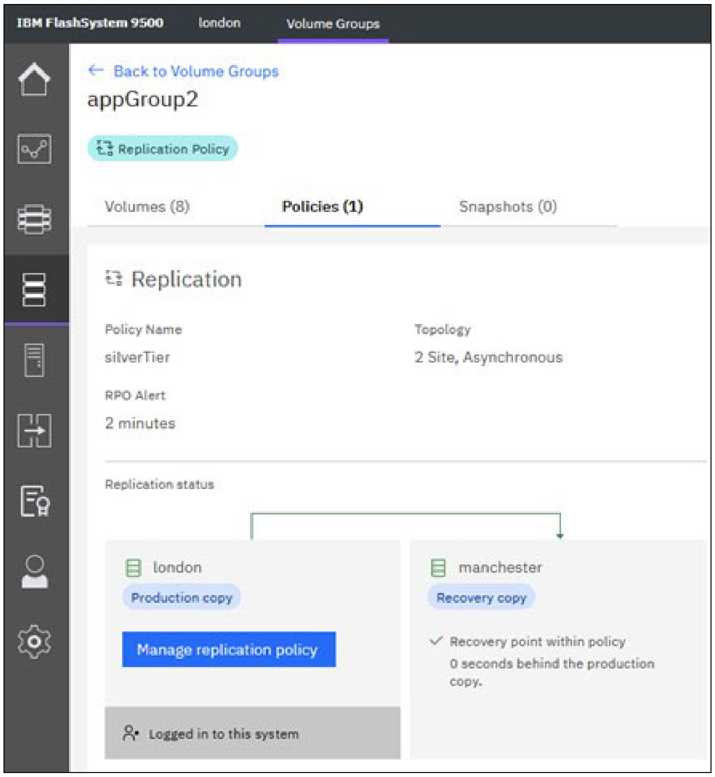

appGroup2

Replication

Policy Name: silverTier

Topology: 2 Site, Asynchronous

RPO Alert: 2 minutes

Replication status

✓ Recovery point within policy
0 seconds behind the production copy.

Figure 4-30 Volume group is updated automatically on the remote system

5. Select the new replication policy to assign to the group. Only eligible policies are listed. A preview of the replication is displayed, which indicates that the direction of replication does not change when the policy is changed.
6. Click **Apply changes** (Figure 4-31).

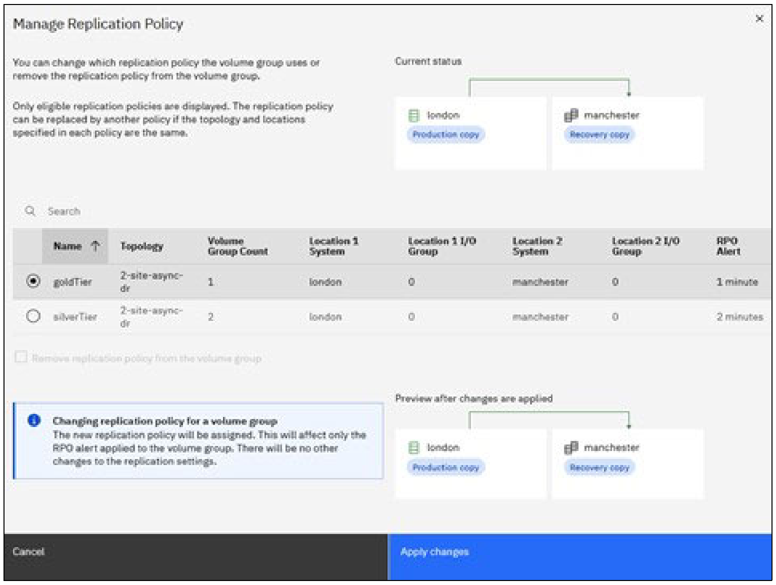
Manage Replication Policy

You can change which replication policy the volume group uses or remove the replication policy from the volume group.

Only eligible replication policies are displayed. The replication policy can be replaced by another policy if the topology and locations specified in each policy are the same.

	Name ↑	Topology	Volume Group Count	Location 1 System	Location 1 I/O Group	Location 2 System	Location 2 I/O Group	RPO Alert
●	goldTier	2-site-async-dr	1	london	0	manchester	0	1 minute
○	silverTier	2-site-async-dr	2	london	0	manchester	0	2 minutes

□ Remove replication policy from the volume group

Changing replication policy for a volume group
The new replication policy will be assigned. This will affect only the RPO alert applied to the volume group. There will be no other changes to the replication settings.

Figure 4-31 Manage Replication Policy

You do not need to repeat the process on the remote system; the volume group is updated automatically on the remote system when the systems are connected.

4.1.13 Unassigning a replication policy from a volume group

Unassigning a replication policy from a volume group unconfigures policy-based replication on the group, which returns the local group to a standalone volume group.

If the replication policy is unassigned from the production copy of a volume group, the recovery copy of the volumes and volume group is automatically removed on the remote system.

To keep the recovery copy of the volumes and volume group as a standalone volume group after the replication policy is removed, you must first enable independent access to the volume group on the recovery system before the replication policy is unassigned.

The GUI provides a preview to allow you to verify the changes before applying them, regardless of whether the systems are disconnected.

Removing the replication policy from the production copy and removing the recovery copy

To remove the replication policy from the production copy, complete these steps using the GUI on the production system:

1. Select **Volumes** → **Volume Groups**.
2. Click on the volume group name to view details for the group.
3. Click the **Policies** tab for the volume group.
4. Click **Manage replication policy**.
5. Select **Remove replication policy from the volume group**.
6. Review the preview to confirm your changes. The preview indicates that the copy of the volumes and volume group on the recovery system will be removed.
7. To proceed with the changes, click **Apply changes** (Figure 4-32).

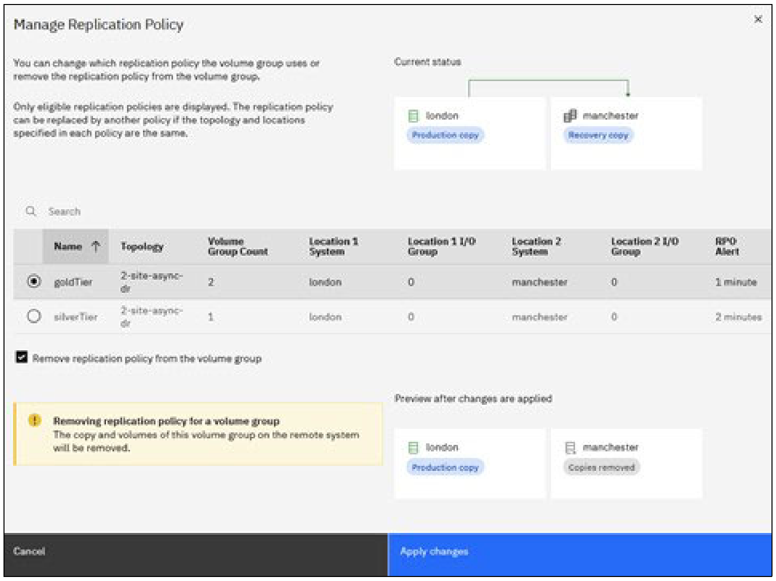
Manage Replication Policy

You can change which replication policy the volume group uses or remove the replication policy from the volume group.

Only eligible replication policies are displayed. The replication policy can be replaced by another policy if the topology and locations specified in each policy are the same.

	Name ↑	Topology	Volume Group Count	Location 1 System	Location 1 I/O Group	Location 2 System	Location 2 I/O Group	RPO Alert
●	goldTier	2-site-async-dr	2	london	0	manchester	0	1 minute
○	silverTier	2-site-async-dr	1	london	0	manchester	0	2 minutes

☑ Remove replication policy from the volume group

Removing replication policy for a volume group
The copy and volumes of this volume group on the remote system will be removed.

Figure 4-32 Manage Replication Policy

Removing the replication policy from an independent copy; keeping both copies

If you want to keep the recovery copy of the volumes after the replication policy is removed, first use the management GUI on the recovery system to enable access to the recovery copy. This makes the two copies of the volume group independent. This section describes how to enabe independent access.

To remove the replication policy from an independent copy of a volume group, complete these steps using the GUI on either system:

1. Select **Volumes** → **Volume Groups**.
2. Click on the volume group name to view details for the group.
3. Click the **Policies** tab for the volume group.
4. Click **Manage replication policy**.
5. Select **Remove replication policy from the volume group**.
6. Review the preview to confirm your changes. The preview indicates that both copies will be retained as standalone groups and the volumes in each copy will continue to be accessible.
7. To proceed with the changes, click **Apply changes** (Figure 4-33).

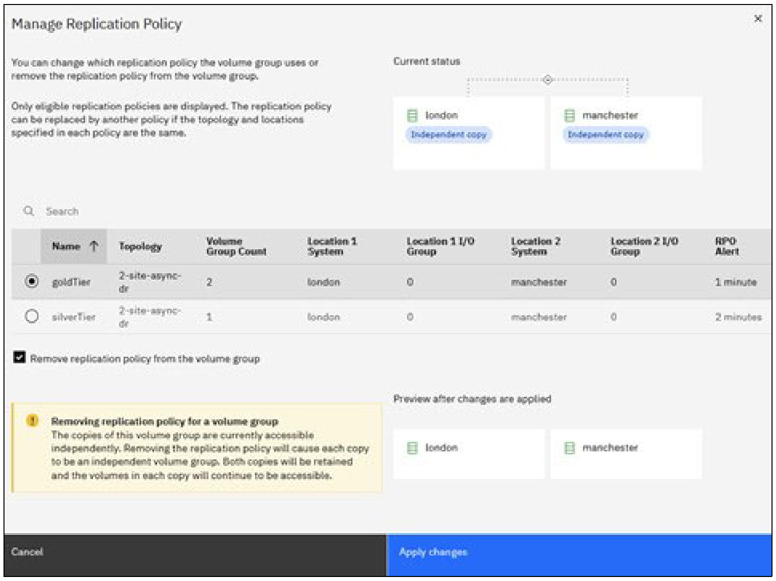

Manage Replication Policy

You can change which replication policy the volume group uses or remove the replication policy from the volume group.

Only eligible replication policies are displayed. The replication policy can be replaced by another policy if the topology and locations specified in each policy are the same.

Current status

london — Independent copy

manchester — Independent copy

Search

	Name ↑	Topology	Volume Group Count	Location 1 System	Location 1 I/O Group	Location 2 System	Location 2 I/O Group	RPO Alert
●	goldTier	2-site-async-dr	2	london	0	manchester	0	1 minute
○	silverTier	2-site-async-dr	1	london	0	manchester	0	2 minutes

☑ Remove replication policy from the volume group

Removing replication policy for a volume group
The copies of this volume group are currently accessible independently. Removing the replication policy will cause each copy to be an independent volume group. Both copies will be retained and the volumes in each copy will continue to be accessible.

Preview after changes are applied

london

manchester

Figure 4-33 Manage Replication Policy

4.1.14 Performing a failover on a volume group (enabling independent access)

To perform a failover to the recovery copy, complete these steps using the GUI on the recovery system:

1. Select **Volumes** → **Volume Groups**.
2. Click on the volume group name to view details for the group.
3. Click the **Policies** tab for the volume group.
4. Review the recovery point that will become accessible.

  - If replication is running, this is displayed as a time that the recovery copy is behind the production copy.
  - If replication is suspended, for example if the systems are disconnected or due to an error, the recovery point is displayed as a fixed timestamp.

5. Click **Enable access (**Figure 4-34).

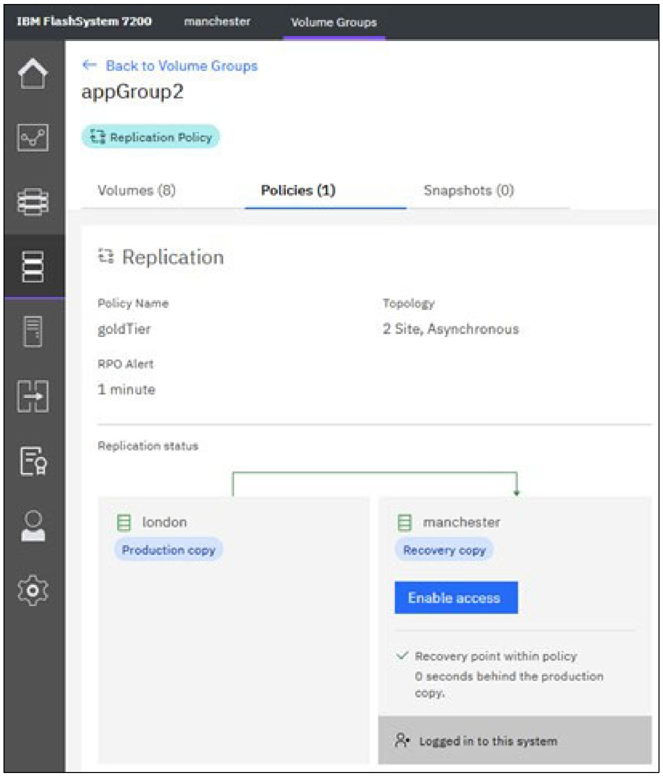

Policy Name
goldTier

Topology
2 Site, Asynchronous

RPO Alert
1 minute

✓ Recovery point within policy
0 seconds behind the production copy.

Logged in to this system

Figure 4-34 Perform a failover to the recovery copy

6. Review the preview to confirm your changes. The preview indicates that replication will be suspended and both copies will be accessible independently for host I/O and configuration changes.
7. To proceed with the changes, click **Enable** (Figure 4-35 on page 50).

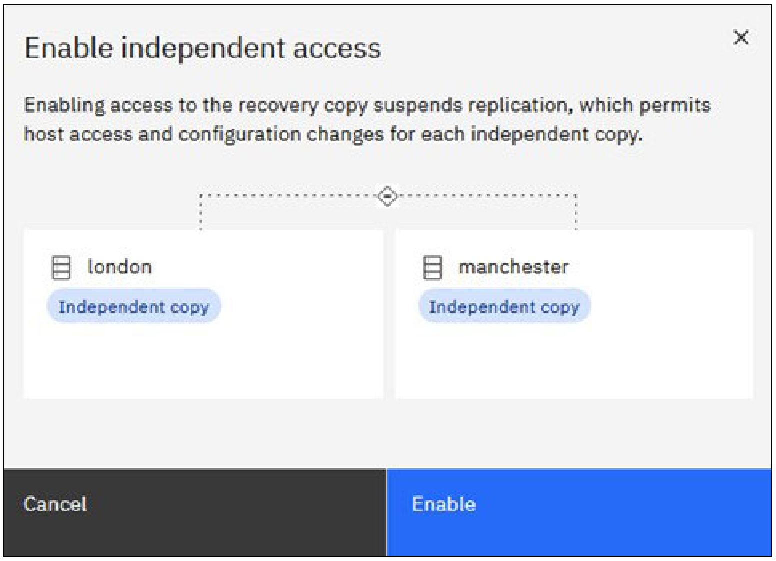
Enable independent access

Enabling access to the recovery copy suspends replication, which permits host access and configuration changes for each independent copy.

london: Independent copy

manchester: Independent copy

Figure 4-35 Enable independent access

Figure 4-36 shows the resulting view of the independent volume groups from the system in location 2 of the policy, which was previously the recovery system.

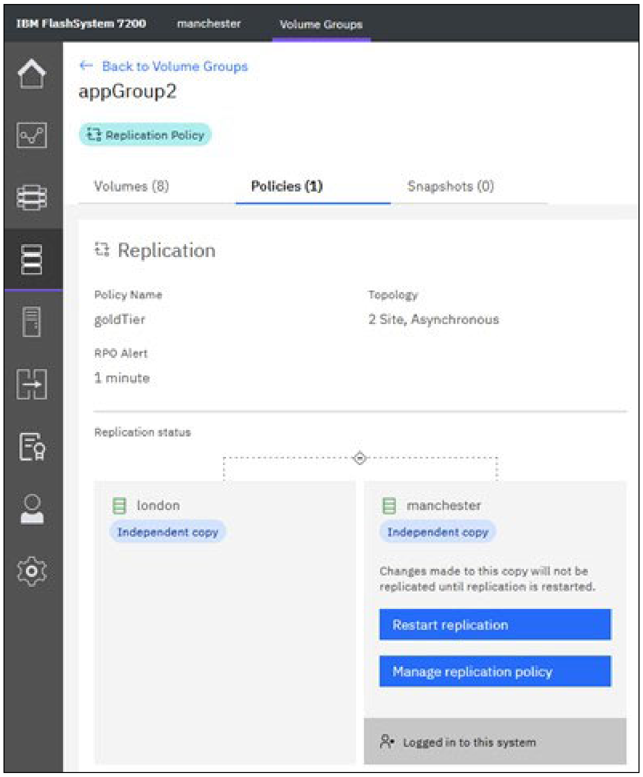
IBM FlashSystem 7200 › manchester › Volume Groups

appGroup2

Replication

Policy Name: goldTier

Topology: 2 Site, Asynchronous

RPO Alert: 1 minute

Replication status

london: Independent copy

manchester: Independent copy

Changes made to this copy will not be replicated until replication is restarted.

Logged in to this system

Figure 4-36 View of the independent volume groups from the system in location 2 of the policy

Similarly, Figure 4-37 on page 51 shows the resulting view of the independent volume groups from the system in location 1 of the policy, which was previously the production system.

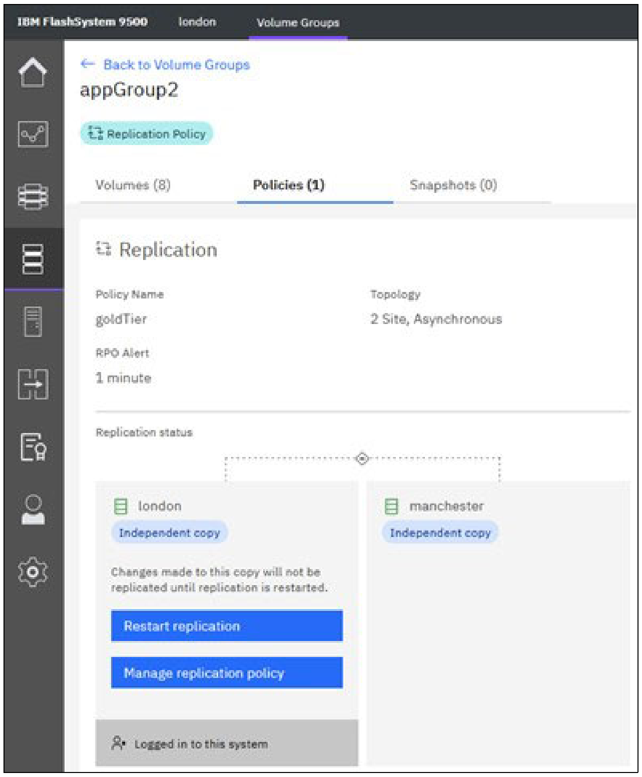

Policy Name: goldTier

Topology: 2 Site, Asynchronous

RPO Alert: 1 minute

Changes made to this copy will not be replicated until replication is restarted.

Figure 4-37 View of the independent volume groups from the system in location 1 of the policy

When a volume group is independent, the volumes in the group are online and accessible for host I/O (Figure 4-38).

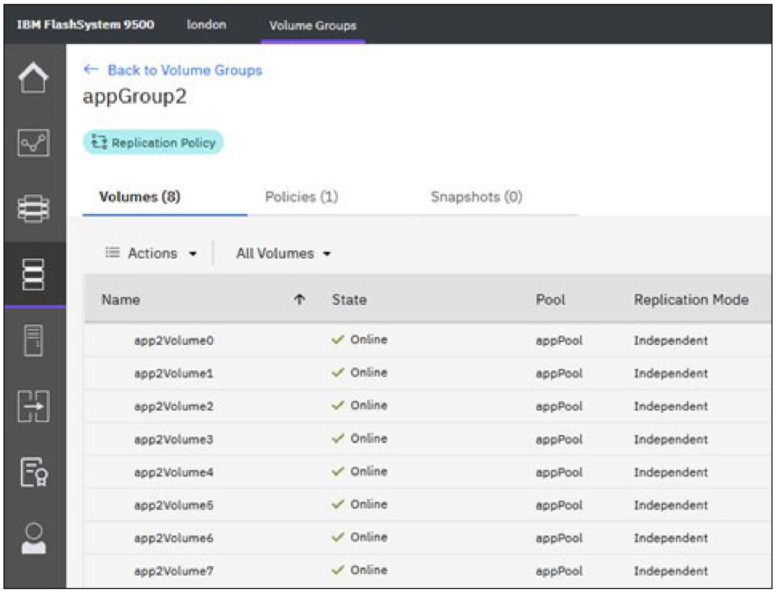

Name	State	Pool	Replication Mode
app2Volume0	✓ Online	appPool	Independent
app2Volume1	✓ Online	appPool	Independent
app2Volume2	✓ Online	appPool	Independent
app2Volume3	✓ Online	appPool	Independent
app2Volume4	✓ Online	appPool	Independent
app2Volume5	✓ Online	appPool	Independent
app2Volume6	✓ Online	appPool	Independent
app2Volume7	✓ Online	appPool	Independent

Figure 4-38 Volumes in the independent group are online

4.1.15 Restarting replication on volume group after a failover

Replication can be restarted by logging into the system that you want to become the production system.

To restart replication, complete these steps using the GUI on the system that includes the copy of the volume group that you want to use as the production copy:

1. Select **Volumes** → **Volume Groups**.
2. Click on the volume group name to view details for the group.
3. Click the **Policies** tab for the volume group.
4. Click **Restart replication**.
5. Review the previous copy direction for the volume group. If host write I/O was performed to the copy after independent access was enabled, this is indicated in the graphic. **Check this information carefully**; any configuration or data that was changed on the remote system will be overwritten if replication is restarted.
6. Review the preview to confirm the changes that will be made if replication is restarted from this system.
7. To proceed with the changes, select **Confirm the volume group on this system should be used as the production copy** and click **Enable** (Figure 4-39).

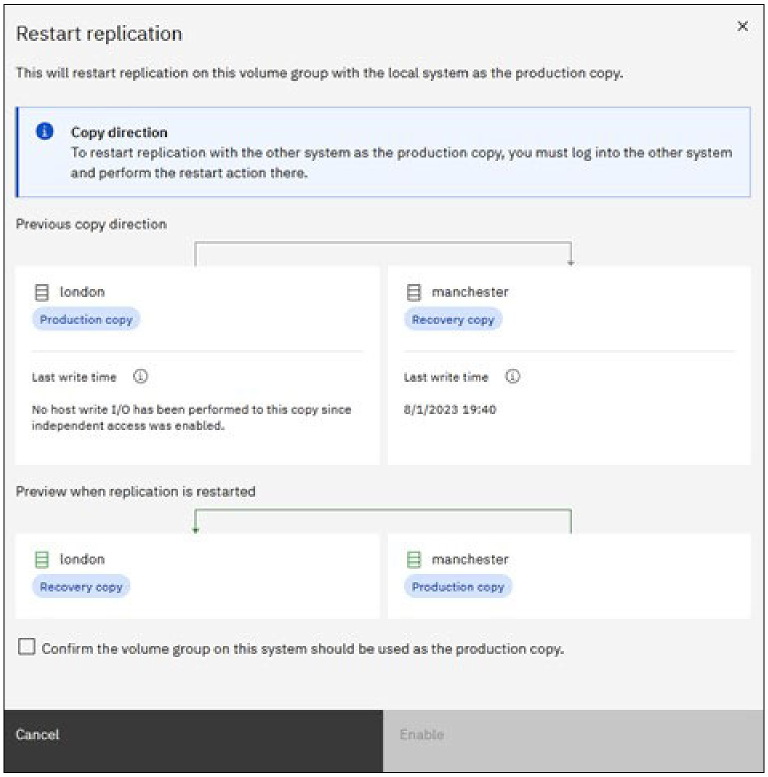

Restart replication

This will restart replication on this volume group with the local system as the production copy.

Copy direction
To restart replication with the other system as the production copy, you must log into the other system and perform the restart action there.

Previous copy direction

london
Production copy

Last write time
No host write I/O has been performed to this copy since independent access was enabled.

manchester
Recovery copy

Last write time
8/1/2023 19:40

Preview when replication is restarted

london
Recovery copy

manchester
Production copy

□ Confirm the volume group on this system should be used as the production copy.

Figure 4-39 Restart replication

The Replication status panel indicates the new status for the volume group. If necessary, the copy progress displays while the recovery copy is being updated with any changes from the production copy (Figure 4-40 on page 53).

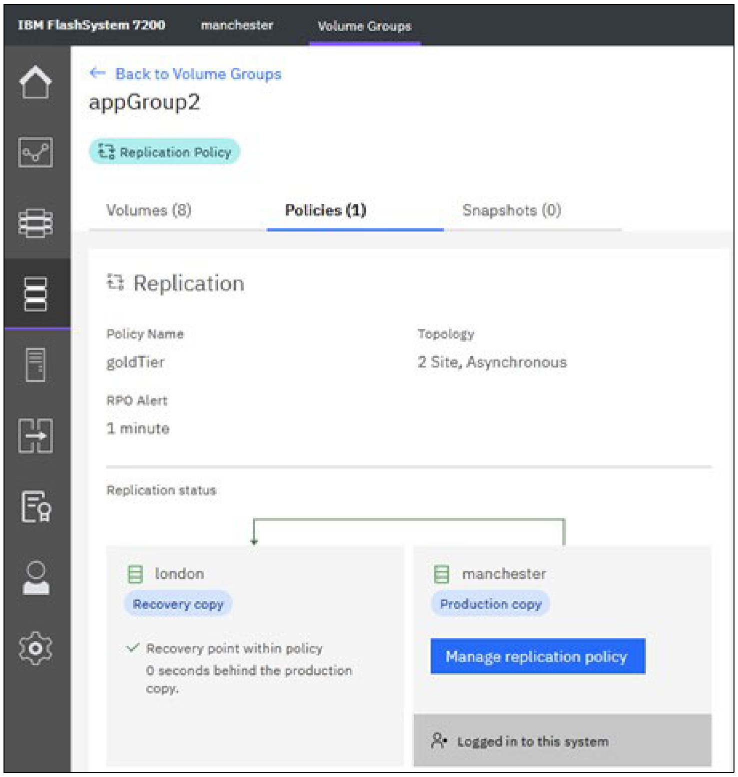

appGroup2

Replication

Policy Name
goldTier

Topology
2 Site, Asynchronous

RPO Alert
1 minute

Replication status

london
Recovery copy

✓ Recovery point within policy
0 seconds behind the production copy.

manchester
Production copy

Figure 4-40 Copy progress displays while the recovery copy is being updated

4.2 Using the CLI

The following section contains examples of how to perform common replication tasks using the CLI. The CLI syntax might change between releases. Ensure that you review the online documentation relating to the version of code that you are using.

4.2.1 Creating partnerships

Before two systems can be configured for replication, they must be in a partnership that can be achieved by using the **`mkfcpartnership`** or **`mkippartnership`** commands.

SSL certificates are exchanged between systems and stored in a trust store. Policy-based replication uses SSL certificates to cross authenticate and allow each system to run commands on the partnered system. This partnership allows most policy-based replication configuration to be done from a single system after the initial setup is complete.

- If configuration is done using the GUI, certificates are generated and exchanged automatically.
- If configuration is done using the CLI, the certificates must be manually generated, transferred, and imported into each system.

4.2.2 SSL certificates (exporting, transferring and the trust store)

On each system, an SSL certificate must be generated and transferred to the system to be partnered. The SSL certificate must then be added to the trust store.

Using the example system names of London and Manchester:

On the London system:

```
svctask chsystemcert -export
```

This command exports the system SSL certificate to a file in */dumps* named *certificate.pem*. Using an SCP client, connect to the system on which you generated the SSL certificate (using the admin user and credentials you used to log in using the CLI). Copy the *certificate.pem* file to your local machine and then use the scp client to copy the *certificate.pem* file to the */tmp* directory on the other system (Manchester in this example).

On the Manchester system:

```
mktruststore -restapi on -file /tmp/certificate.pem -name london
```

This command adds the SSL generated on the other (London) system to the truststore on the Manchester System. Naming is optional but with multiple partnerships it helps identify which certificates are in use.

You can check the contents of the Trust Store and see the imported certificate using the **lstruststore** command:

```
IBM_FlashSystem:manchester:superuser>lstruststore

id name    percent_used  space    restapi   ipsec   vasa
0  london  10            10.7KB   on        off     off
```

Repeat the steps above, but export the system certificate from the Manchester system and import it into the London system.

When this process is complete, both systems have exported SSL certificates, which have been transferred to the systems to be partnered, and imported into the truststores.

4.2.3 Fibre channel partnership

The syntax for creating a FC partnership is as follows:

```
mkfcpartnership -linkbandwidthmbits <link_bandwidth_in_mbps> [-backgroundcopyrate
<percentage>] remote_system_id | remote_system_name
```

Example 4-1 On London

```
mkfcpartnership -linkbandwidthmbits 16000 -backgroundcopyrate 100 Manchester
```

Example 4-2 On Manchester

```
mkfcpartnership -linkbandwidthmbits 16000 -backgroundcopyrate 100 London
```

4.2.4 IP partnership

The syntax for creating an IP partnership is as follows:

```
mkippartnership -type ipv4|ipv6 -clusterip <ipadr> [-chapsecret <CHAPsecret>]
-linkbandwidthmbits <link_bandwidth_in_mbps> -backgroundcopyrate <percentage>
-compressed yes|no remote_system_id | remote_system_name
```

Example 4-3 On London

```
mkippartnership -type ipv4 -clusterip 192.168.32.19 -chapsecret mychapsecret
-linkbandwidthmbits 100 -backgroundcopyrate 100 -compressed yes Manchester
```

Example 4-4 On Manchester

```
mkippartnership -type ipv4 -clusterip 192.168.32.29 -chapsecret mychapsecret
-linkbandwidthmbits 100 -backgroundcopyrate 100 -compressed yes London
```

4.2.5 Creating Replication Policies

A *Replication Policy* can be defined between two policy-based replication-capable systems that are in a fully-configured partnership.

The system names can be displayed using the **`lspartnership`** command. Currently, the only supported topology is `2-site-async-dr`. The `RPO alerting threshold` is expressed in seconds.

Each replication policy identifies a single I/O group on each system; if multiple I/O groups are required for replication in each system, then multiple replication policies must be created.

Figure 4-41 shows the syntax for creating a replication policy.

```
>>- mkreplicationpolicy --+------------------------------------+-->
                          '- -name -- replication_policy_name--'

>-- -topology -- 2-site-async-dr --------------------------------->

>-- -location1system --+- location1_system_name -+---------------->
                       '- location1_system_id ---'

>-- -location1iogrp -- location1_iogrp_id ------------------------>

>-- -location2system --+- location2_system_name -+---------------->
                       '- location2_system_id ---'

>-- -location2iogrp -- location2_iogrp_id ------------------------>

>-- -rpoalert -- rpo_alert_threshold_in_seconds -----------------><
```

Figure 4-41 Syntax for creating a replication policy

Example 4-5 Creating replication policy example

```
svctask mkreplicationpolicy -name ldn_to_mcr_policy -topology 2-site-async-dr
-location1system lodon -location1iogrp 0 -location2system manchester
-location2iogrp 0 -rpoalert 300
```

4.2.6 Creating provisioning policies and associating with pools

An optional provisioning policy allows an administrator to simplify the volume provisioning process by specifying the capacity savings on a 'per pool' basis, in addition to any provided by the back-end media. A provisioning policy can be associated with multiple pools, and defines whether volumes that are created within are compressed, thin, or none (uncompressed unless using FCMs in which case only the native FCM compression will be used).

Figure 4-42 shows the syntax for creating a provisioning policy.

```
>>- mkprovisioningpolicy -- -+-------------+----------------->
                             '- -name name -'

>-- -capacitysaving --+- none -------+--+----------------+--><
                      +- thin -------+  '- -deduplicated -'
                      '- compressed -'
```

Figure 4-42 Syntax for creating a provisioning policy

Example 4-6 Creating a provisioning policy example

```
mkprovisioningpolicy -name london_thin -capacitysaving thin
```

Multiple provisioning policies can be defined, but each Pool, or child Pool, can only have a single Provisioning Policy assigned.

After a provisioning policy is defined, it can be associated with a Pool when the Pool is created, or an existing Pool can be modified to add the association.

Example 4-7 Associating a provisioning policy with a pool

```
mkmdiskgrp -ext 1024 -provisioningpolicy london_thin
chmdiskgrp -provisioningpolicy london_thin Pool0
```

If a Pool no longer requires a provisioning policy, the association can be removed.

Example 4-8 Removing the association

```
chmdiskgrp -noprovisioningpolicy Pool0
```

4.2.7 Creating child pools

Creating child pools allows the user to have flexibility as each storage pool can only have one Pool Link, and a single Provisioning Policy assigned. Child Pools can be created within a storage pool and each child pool can also have its own Pool Link and Provisioning Policy assigned to participate in Policy Based Replication.

For example, after creating a data reduction pool, *mdiskgrp0*, child pools can be created as follows:

```
mkmdiskgrp -parentmdiskgrp mdiskgrp0 -datareduction yes -name pbr_child1 -noquota
mkmdiskgrp -parentmdiskgrp mdiskgrp0 -datareduction yes -name pbr_child2 -noquota
```

4.2.8 Creating pool links

Normally, we expect pool links to be created using the GUI. However, it is possible to do this from the CLI by manually setting both storage pools to have the same pool link ID.

One way to do this is to record the *replication_pool_link_uid* from the pool on the DR cluster, and use this to set the *replication_pool_link_uid* on the Production cluster. The objective is for both systems to have the same *replication_pool_link_uid* for the linked pools.

On the DR cluster, use the following command:

```
lsmdiskgrp dr_pool | grep replication_pool_link_uid
```

Example 4-9 Example output

```
lsmdiskgrp dr_pool | grep replication_pool_link_uid
replication_pool_link_uid 00000000000000010000020431C02050
```

On the Production cluster, use the following command:

```
chmdiskgrp -replicationpoollinkuid 00000000000000010000020431C02050 prod_pool
```

> **Note:** The previous example assumes that the Storage Pool on the DR cluster is named *dr_pool*, and the pool to be linked on the production cluster is named *prod_pool*.

4.2.9 Creating volume groups with a Replication Policy

In order to replicate using policy-based replication, volumes need to be a member of a volume group, which has an associated Replication Policy.

The Volume Group can have the Replication Policy associated on creation, as is described here, or it can be added afterward, as is described in 4.2.10, “Associating an existing volume group with a replication policy” on page 57.

The replication policy to be added can be specified by name, or by its ID.

Example 4-10 Creating volume groups with a Replication Policy

```
mkvolumegroup -replicationpolicy two_site_policy -name copy_me
```

4.2.10 Associating an existing volume group with a replication policy

A Volume Group without an assigned replication policy can have a replication policy assigned by using the **chvolumegroup** command. This command can also be used to change the replication policy that is assigned to a Volume Group, or to remove the replication policy if the contents of the Volume Group are no longer required to be replicated.

Figure 4-43 shows the **chvolumegroup** command syntax.

```
>>- chvolumegroup ---------------------------------------------------->

>--+---------------------------------------------------------------+-->
   +- -replicationpolicy --+- replication_policy_name -+-----------+
   |                       '- replication_policy_id ---'           |
   '- -noreplicationpolicy ----------------------------------------'

>--+- volumegroup_id ---+--------------------------------------------><
   '- volumegroup_name -'
```

Figure 4-43 chvolumegroup command syntax

Example 4-11 Assign a replication policy to a Volume Group named production_group

```
chvolumegroup -replicationpolicy hampshire_policy production_group
```

4.2.11 Adding volumes to volume groups

Volumes can be added into a volume group using the **chvdisk** command.

Example 4-12 Adding volumes to volume groups

```
chvdisk -volumegroup ldn_manc_group database_vol_01
```

Alternatively, new volumes can have the volume group specified when they are created.

Example 4-13 Specifying the volume group when the volume is created

```
mkvolume -size 500 -unit gb -pool replicated_pool -volumegroup ldn_manc_group
 -name database_vol_01
```

4.2.12 Viewing information about the volume group

To see the state of a volume group that is engaged in replication, with a replication and provisioning policy assigned, the **lsvolumegroupreplication** command can be used.

The **lsvolumegroupreplication** command summarizes the volume groups. If a replication policy is associated, it is shown in the *replication_policy_name* column.

```
IBM_FlashSystem:gobibear:superuser>lsvolumegroupreplication
id    name       replication_policy_id   replication_policy_name
0     copy_me    0                       two_site_policy
```

The detailed view shows the current state of a replicated volume group.

```
IBM_FlashSystem:gobibear:superuser>lsvolumegroupreplication 0
id 0
name copy_me
replication_policy_id 0
replication_policy_name two_site_policy
location1_system_id 000002042F4101DE
location1_system_name gobibear
location1_replication_mode production
location1_status healthy
location1_running_recovery_point
location1_fixed_recovery_point
location1_within_rpo
location1_volumegroup_id 0
location1_sync_required
location1_sync_remaining
location1_previous_replication_mode
location1_last_write_time
location2_system_id 000002042E410222
location2_system_name cavebear
location2_replication_mode recovery
location2_status healthy
location2_running_recovery_point 0
location2_fixed_recovery_point
```

```
location2_within_rpo yes
location2_volumegroup_id 0
location2_sync_required
location2_sync_remaining
location2_previous_replication_mode
location2_last_write_time
link1_status running
```

Key fields are the *location1_status* and *location2_status* that show whether replication is running and healthy. The l*ocation2_within_rpo* line shows that the DR site is within the RPO that is as part of the replication policy.

The next example shows an example of a volume group where replication stopped (because the DR cluster is offline):

```
IBM_FlashSystem:gobibear:superuser>lsvolumegroupreplication 0
id 0
name copy_me
replication_policy_id 0
replication_policy_name two_site_policy
location1_system_id 000002042F4101DE
location1_system_name gobibear
location1_replication_mode production
location1_status healthy
location1_running_recovery_point
location1_fixed_recovery_point
location1_within_rpo
location1_volumegroup_id 0
location1_sync_required
location1_sync_remaining
location1_previous_replication_mode
location1_last_write_time
location2_system_id 000002042E410222
location2_system_name cavebear
location2_replication_mode disconnected
location2_status disconnected
location2_running_recovery_point
location2_fixed_recovery_point
location2_within_rpo
location2_volumegroup_id 0
location2_sync_required
location2_sync_remaining
location2_previous_replication_mode
location2_last_write_time
link1_status disconnected
```

4.2.13 DR failover (independent access)

To suspended replication and allow the DR copies of the data to be accessed from a host, use the **chvolumegroupreplication -mode independent** command.

Figure 4-44 on page 60 shows the **chvolumegroupreplication** command syntax.

```
>>- chvolumegroupreplication ----------------------------------------->

>--+-----------------------------+----------------------------------->
   '- -mode --+- independent ---+-'
              '- production ----'

>--+- volumegroup_id ---<
   '- volumegroup_name -'
```

Figure 4-44 chvolumegroupreplication command syntax

This command must be run on the system that is hosting the recovery copy of the volume group. The command will fail if it is attempted on the system that is hosting the production copy of the volume group.

If the command is successful, the state of the volume group can be shown using the **`lsvolumegroupreplication`** command. This includes the state of the volume group, which is now independent.

While the volume group is in *independent* mode, the volumes can be accessed from a host. They can be mapped to a host at any time, but while the replication mode is *recovery* the volumes appear offline on the host and are not available for reading or writing to until the replication mode is changed to *`independent`*.

To restart replication, use the **`chvolumegroupreplication -mode production`** command on the production system.

4.2.14 Deleting volumes

To delete a replicated volume, use the **`rmvolume`** command on the production system. The volume is removed from both the production and the DR systems.

4.2.15 Removing a replication policy

To remove the replication policy associated with a volume group, use the **`chvolumegroup -noreplicationpolicy`** command. This command stops replication of the volumes in the volume group and the DR copies of the volumes are deleted. If you want to keep the DR copies while stopping replication, you must first set the volume group on the DR system to *independent*.

Example 4-14 Removing a replication policy

```
IBM_FlashSystem:jazzy-c:superuser>chvolumegroup -noreplicationpolicy
ldm_manc_group
```

4.2.16 Changing the replication policy of a volume group

A replication policy is immutable and cannot be modified; it can only be created and deleted. If a change is required to the replication policy, a new replication policy must be created with the **`mkreplicationpolicy`** command, and assigned to the volume groups where the updated replication policy is required.

The replication policy that is associated with a volume group can be replaced by another policy if the topology and locations that are specified in each policy are the same.

Example 4-15 Changing the replication policy associated with a volume group

```
svctask chvolumegroup -replicationpolicy twoSiteAsynchronous
replicated_volume_group_1
```

4.2.17 Deleting volume groups

A volume group can be deleted only when it does not contain volumes. To delete a volume group, use the **`chvdisk`** command and set each volume in turn to have *novolumegroup*.

Example 4-16 Deleting volume groups

```
chvdisk -novolumegroup production_vol_1
rmvolumegroup production_vg
```

4.3 Using the REST API

The process of configuring, managing, and monitoring policy-based replication using the REST API is similar to using the CLI. For details on how to use the REST API in general, consult the documentation.

4.4 Command-line interface quick reference

In this section, several command line options for policy based replication are described.

4.4.1 Provisioning policy commands

mkprovisioningpolicy	Creates a provisioning policy
chprovisioningpolicy	Modifies a provisioning policy
lsprovisioningpolicy	Displays the provisioning policies
rmprovisioningpolicy	Removes a provisioning policy

Using provisioning policies with storage pools

To create a storage pool with a provisioning policy:

```
mkmdiskgrp -provisioningpolicy <policy_id/name> ...
```

To assign a provisioning policy to a storage pool:

```
chmdiskgrp -provisioningpolicy <policy_id/name> <pool_id/name>
```

To display the provisioning policy (*provisioning_policy_id*, *provisioning_policy_name*) assigned to storage pools:

```
lsmdiskgrp
```

To unassign a provisioning policy from a storage pool:

```
chmdiskgrp -noprovisioningpolicy <pool_id/name>
```

4.4.2 Pool linking commands

To display the replication pool link unique identifier (*replication_pool_link_uid*) assigned to storage pools:

```
lsmdiskgrp
```

To link a pool, set the replication pool link identifier on the pool, linking it with pools on partnered systems that have the same value:

```
chmdiskgrp -replicationpoollinkuid <uid> <pool_id/name>
```

To unlink a pool, reset the replication pool link identifier on the pool:

```
chmdiskgrp -resetreplicationpoollinkuid <pool_id/name>
```

4.4.3 Replication policy commands

mkreplicationpolicy	Creates a replication policy
lsreplicationpolicy	Displays the replication policies
rmprovisioningpolicy	Removes a replication policy

Using replication policies with volume groups

To create a volume group with a replication policy:

```
mkvolumegroup -replicationpolicy <policy_id/name> ...
```

To assign a replication policy to a volume group:

```
chvolumegroup -replicationpolicy <policy_id/name> <volumegroup_id/name>
```

To display the replication policy (*replication_policy_id*, *replication_policy_name*) assigned to volume groups:

```
lsvolumegroup
```

To unassign a replication policy from a volume group:

```
chvolumegroup -noreplicationpolicy <volumegroup_id/name>
```

4.4.4 Volume group replication commands

lsvolumegroupreplication	Displays the volume group replication information
chvolumegroupreplication	Modifies the replication settings of a volume group

To change a volume group from a recovery copy to an independent copy, enabling access to the volumes in the group on the recovery system:

```
chvolumegroupreplication -mode independent <volumegroup_id/name>
```

The system automatically changes the production copy on the remote system to an independent copy when the systems are connected.

To change a volume group from an independent copy to a production copy, restarting replication with the local system as the production copy:

```
chvolumegroupreplication -mode production <volumegroup_id/name>
```

The system automatically changes the remote volume group to a recovery copy.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only.

- *Performance and Best Practices Guide for IBM Spectrum Virtualize 8.5,* SG24-8521
- *Implementation Guide for IBM Spectrum Virtualize Version 8.5*, SG24-8520
- *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574-02

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- Policy-based replication

  https://www.ibm.com/docs/en/flashsystem-7x00/8.5.x?topic=replication-policy-based

- IBM FlashSystem 5x00 documentation

  https://www.ibm.com/docs/en/flashsystem-5x00

- IBM FlashSystem 7x00 documentation

  https://www.ibm.com/docs/en/flashsystem-7x00

- IBM FlashSystem 9x00 documentation

  https://www.ibm.com/docs/en/flashsystem-9x00

- IBM SAN Volume Controller (2145 and 2147) documentation

  https://www.ibm.com/docs/en/sanvolumecontroller

- IBM Spectrum Virtualize for Public Cloud documentation

  https://www.ibm.com/docs/en/spectrumvirtualizecl

Help from IBM

- IBM Support and downloads

  ibm.com/support

- IBM Global Services

  ibm.com/services

REDP-5704-00

ISBN 0738461032

Printed in U.S.A.