IBM

IBM SAN Volume Controller and Storwize Family Native IP Replication

- Implement SVC/Storwize Family native IP replication
- Understand requirements for configuring native IP replication
- Mitigate the effects of IP latency over distance

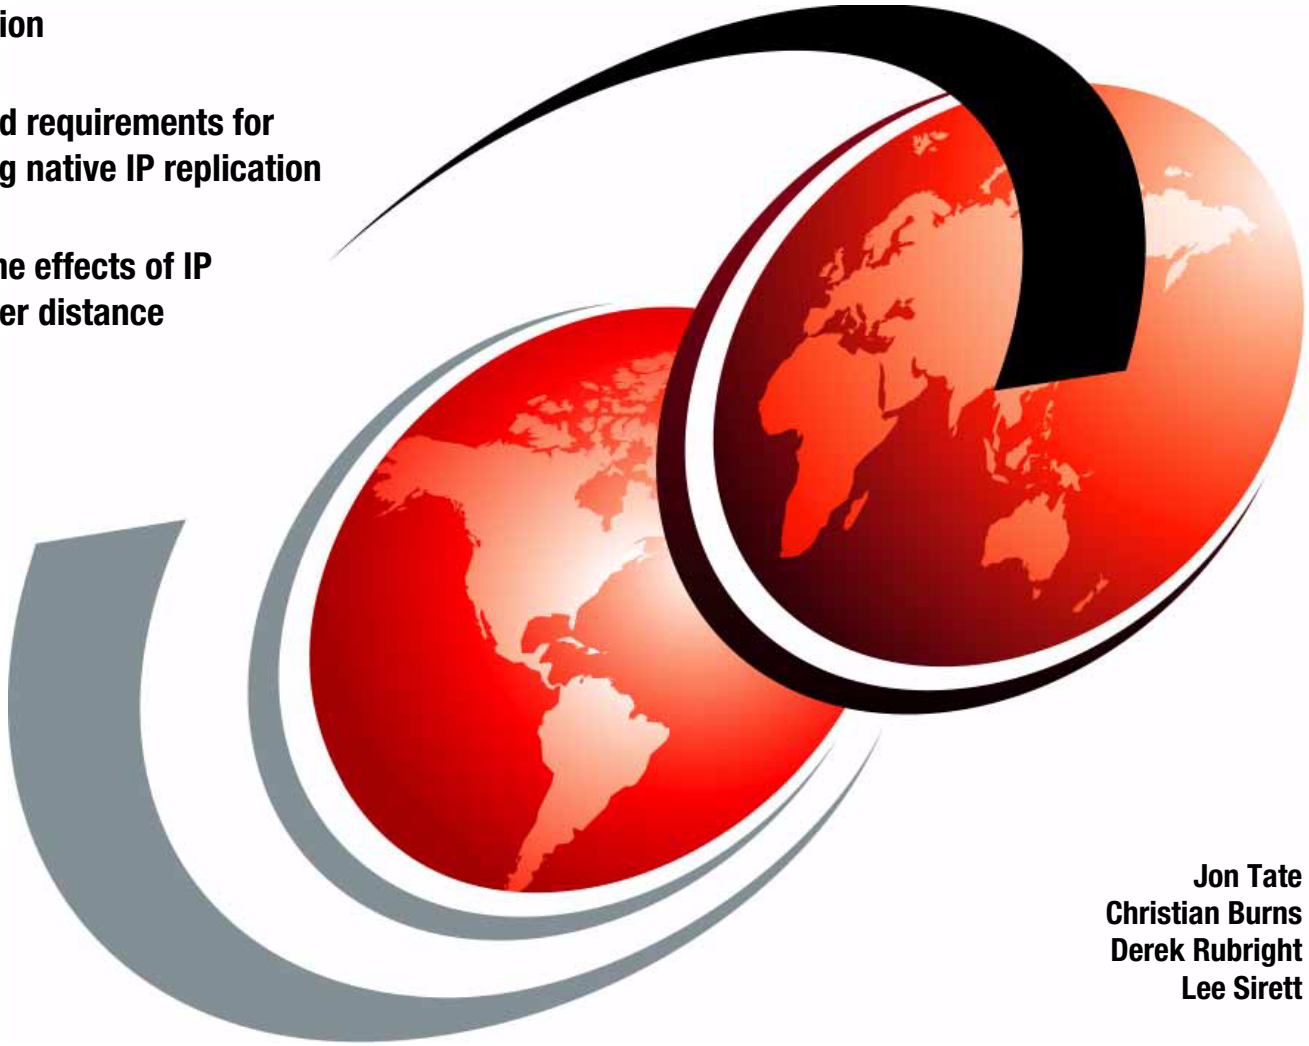

Jon Tate
Christian Burns
Derek Rubright
Lee Sirett

Redpaper

ibm.com/redbooks

International Technical Support Organization

IBM SAN Volume Controller and Storwize Family Native IP Replication

September 2014

REDP-5103-00

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

First Edition (September 2014)

This edition applies to version 7.2 and above of IBM SAN Volume Controller and Storwize V7000, Storwize V5000, and Storwize V3700.

Contents

Notices

This information was developed for products and services offered in the U.S.A.

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, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 states 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 supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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 the names and addresses used by an actual business enterprise 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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

DB2®
FlashCopy®
FlashSystem™
IBM®
IBM FlashSystem™
Power Systems™
Real-time Compression™
Redbooks®
Redpaper™
Redbooks (logo) ®
Storwize®
System Storage®
XIV®

The following terms are trademarks of other companies:

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

Preface

IBM® has announced native Internet Protocol (IP) replication using Bridgeworks SANSlide technology with its IBM System Storage® SAN Volume Controller (SVC), IBM Storwize® V7000, IBM Storwize V5000 and Storwize V3700 virtualized storage systems.

This combination of SANSlide and the SVC/Storwize family provides a powerful solution for clients who require efficient, IP-based replication over long distances.

This certification gives SVC/Storwize clients a fully supported, transparent technology that includes unmatched levels of performance and reliability. With the SANSlide protocol acceleration technology, it is now possible to replicate data across continents in a cost-efficient way, with little or no loss in performance. At the same time, bandwidth usage can improve to over 95%, rather than the 1% - 5% normally achieved in long-distance IP networks.

This IBM Redpaper™ publication shows the steps required to implement this solution efficiently and speedily.

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Hursley Lab, UK.

Jon Tate is a Project Manager for IBM System Storage SAN Solutions at the ITSO, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2 support for IBM storage products. Jon has 28 years of experience in storage software and management, services, and support, and is both an IBM Certified IT Specialist and an IBM SAN Certified Specialist. He is also the UK Chairman of the Storage Networking Industry Association.

Christian Burns is an IBM Storage Solution Architect based in New Jersey. As a member of the Storage Solutions Engineering team in Littleton, MA, he works with clients, IBM Business Partners, and IBMers worldwide, designing and implementing storage solutions that include various IBM products and technologies. Christian's areas of expertise include IBM Real-time Compression™, SVC, IBM XIV®, and IBM FlashSystem™. Before joining IBM, Christian was the Director of Sales Engineering at IBM Storwize, before it became IBM Storwize. He brings over a decade of industry experience in the areas of sales engineering, solution design, and software development. Christian holds a BA degree in Physics and Computer Science from Rutgers College.

Derek Rubright is a Staff Software Engineer, specializing in IBM SAN Volume Controller and Storwize V7000, at Research Triangle Park in Raleigh, NC. He has over 15 years of experience in technical customer support, and 5 years of experience in storage and storage area network (SAN) environments.

Lee Sirett is a Storage Technical Advisor for the European Storage Competency Centre (ESCC) in Mainz, Germany. Before joining the ESCC, he worked in IBM Technical Support Services for 10 years providing support on a range of IBM Products, including IBM Power Systems™. Lee has 24 years of experience in the IT Industry. He is IBM Storage Certified, and both an IBM Certified XIV Administrator and Certified XIV Specialist.

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

Bill Passingham
Andrew Martin
Andrea Sipka
Matt Smith
IBM Hursley

Aakanksha Mathur
Pravin Shinde
Shrirang Bhagwat
IBM India

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 ITSO 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.

Obtain more information about the residency program, browse the residency index, and apply online:

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:

  ibm.com/redbooks

- Send your comments in an email:

  redbooks@us.ibm.com

- Mail your comments:

  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for 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 to native IP replication in SVC/Storwize

This chapter describes the new functionality introduced with code version 7.2, which supports native IP replication across the IBM Storwize family, including IBM System Storage SAN Volume Controller (SVC). The functionality of the new feature is covered, along with a summary of the SVC/Storwize remote mirror technologies.

Important: Although this paper provides information about the Storwize V7000 and SVC, it also applies to the Storwize V5000 and the Storwize V3700.

1.1 IP replication overview

The native Internet Protocol (IP) replication feature enables replication between any SVC/Storwize family member running code version 7.2 or higher, using the built-in networking ports of the cluster nodes. Following a recent partnership with IBM, it uses SANslide technology developed by Bridgeworks Limited of Christchurch, UK. They specialize in products that can bridge storage protocols and accelerate data transfer over long distances.

Adding this technology at each end of a wide area network (WAN) Transmission Control Protocol/Internet Protocol (TCP/IP) link significantly improves the usage of the link. It does this by applying patented Artificial Intelligence (AI) to hide latency normally associated with WANs. Therefore, doing so can greatly improve the performance of mirroring services, in particular Global Mirror with Change Volumes (GM/CV) over long distances.

More details are provided in 1.4, “SVC/Storwize IP replication in more detail” on page 7.

Tip: Although all versions of 7.2 code offer native IP replication, versions 7.2.0.4 and later offer performance enhancements to the SANslide technology.

1.2 Metro Mirror and Global Mirror overview

Metro Mirror (MM) and Global Mirror (GM) are technologies that enable you to keep a real time, or near real-time, copy of a disk at a remote site that contains another SVC cluster or Storwize V7000 system.

1.2.1 Metro Mirror

MM makes *synchronous* copies, which means that the original writes are not considered complete until the write to the destination disk has been confirmed. The distance between two sites is usually determined by how much latency the applications can handle. Therefore, MM is typically used within metropolitan distances in conjunction with a zero recovery point objective (RPO), that is, zero data loss.

Figure 1-1 shows the order of MM write operations.

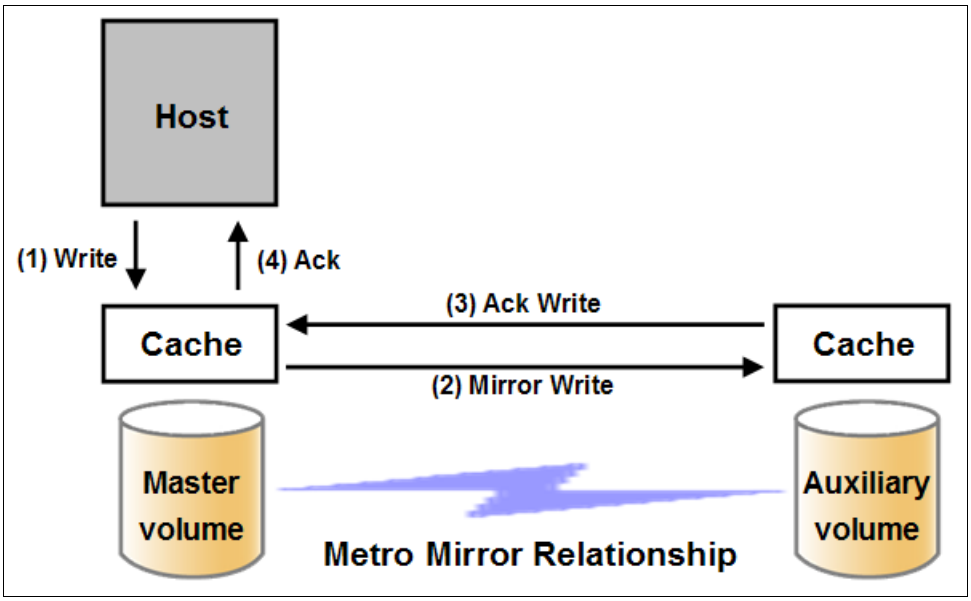

Figure 1-1 Metro Mirror write sequence

1.2.2 Global Mirror

GM makes *asynchronous* copies of your disk. This means that the write is considered complete after it is complete at the local disk; it does not wait for the write to be confirmed at the remote cluster like MM does. This greatly reduces the latency experienced by applications if the other cluster is far away. However, it also means that, during a failure, the data on the remote copy might not have the most recent changes committed to the local disk.

Figure 1-2 shows the order of GM write operations.

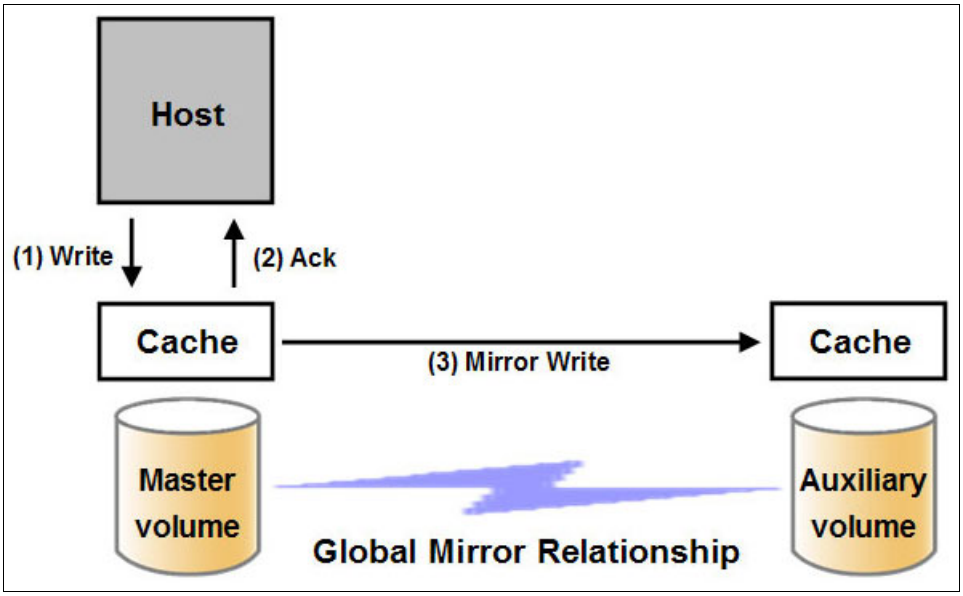

Figure 1-2 Global Mirror write sequence

1.2.3 Global Mirror with Change Volumes

This function (also known as *Cycle-Mode Global Mirror*), was introduced in Storwize V7000 V6.3, and can best be described as *Continuous Remote IBM FlashCopy®*. If you use this feature, the Storwize V7000 will essentially take periodic flash copies of a disk, and write them to your remote destination.

This feature completely isolates the local copy from WAN issues, such as line glitches, and from sudden spikes in workload that might occur. Bear in mind that your remote copy might lag behind the original, depending on how you have set up the cycle time.

Figure 1-3 shows a high-level conceptual view of GM/CV. GM/CV uses FlashCopy to maintain image consistency, and to isolate host volumes from the replication process.

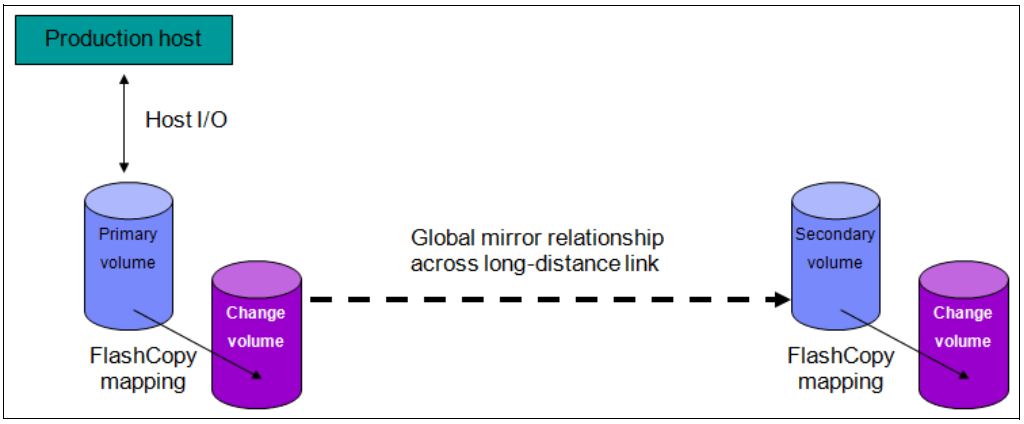

Figure 1-3 Global Mirror with Change Volumes

Information: For more information about SVC and Storwize V7000 replication, see *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574.

For more information about the SVC, see *Implementing the IBM System Storage SAN Volume Controller V7.2*, SG24-7933.

For more information about the Storwize V7000, see *Implementing the IBM Storwize V7000 V7.2*, SG24-7938.

For more information about the Storwize V5000, see *Implementing the IBM Storwize V5000*, SG24-8162

For more information about the Storwize V3700, see *Implementing the IBM Storwize V3700*, SG24-8107

1.3 Example scenarios

A SVC/Storwize cluster in one location can be connected to another SVC/Storwize cluster at a remote location using various designs. The following sections show example connectivity designs.

1.3.1 Single WAN link to a single input/output group

Figure 1-4 shows a single WAN link connecting two Storwize V7000 clusters consisting of a single input/output (I/O) group each. In this design, there is a redundant connection from each Storwize V7000. Either Ethernet port 1 or 2 can be configured for IP replication, but not both. This protects against a node connection failure or a node canister outage (for example, for a code upgrade). However, it does not protect against a WAN connection failure.

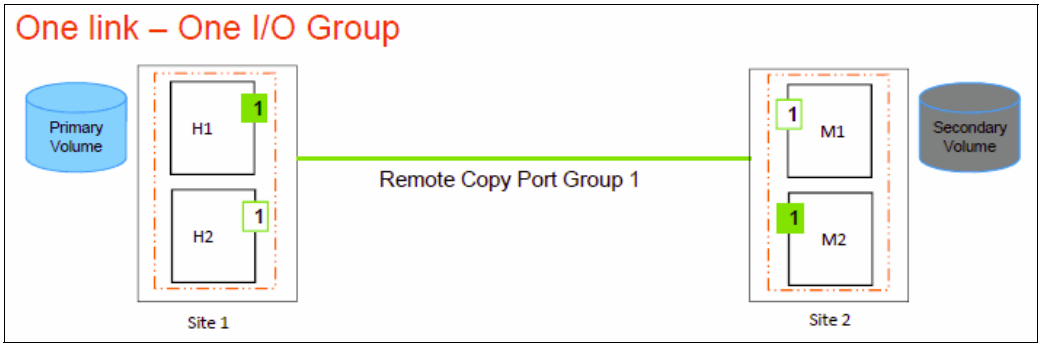

Figure 1-4 Single WAN link to a single I/O group

In this scenario, the remote-copy port group on each system includes two IP addresses, one from each node. When the system initially configures, one of these IP addresses is chosen from each side. This pairing cannot be chosen or changed by an administrator. In the previous example, H1 and M2 have established a connection. If M2 fails, the session will be lost but the connection will automatically be re-established between either H1 or H2 and M1.

The IP addresses do not fail over between nodes for the purposes of IP replication. Although the address does in fact move to the other node, as can be seen from the output of the **`lsportip`** command when a node does fail, this is for Internet Small Computer System Interface (iSCSI) purposes only. The IP replication link will be re-established with the previously assigned original IP address of the port.

1.3.2 Single WAN link to dual I/O groups

Figure 1-5 on page 6 shows a single WAN link connecting two Storwize V7000 clusters consisting of dual I/O groups. This design also provides redundancy against a node failure, but not against a WAN failure. The remote-copy port group on each system has four IP addresses, one from each node.

When the system initially configures, one of these IP addresses is chosen from each side. This pairing cannot be chosen or changed by an administrator. If one of the paired nodes fails, for example H1, then the connection will be lost as before, but re-established between either H2, H3, or H4 and M1, M2, M3, or M4.

Restriction: Storwize V7000 and SVC nodes with more than two I/O groups are supported. However, there is a limitation that only four IP addresses, from two I/O groups per system, can be included in a remote-copy port group. If there is only one intersite link there must be only one port group.

Figure 1-5 shows a single WAN link connecting two Storwize V7000 clusters.

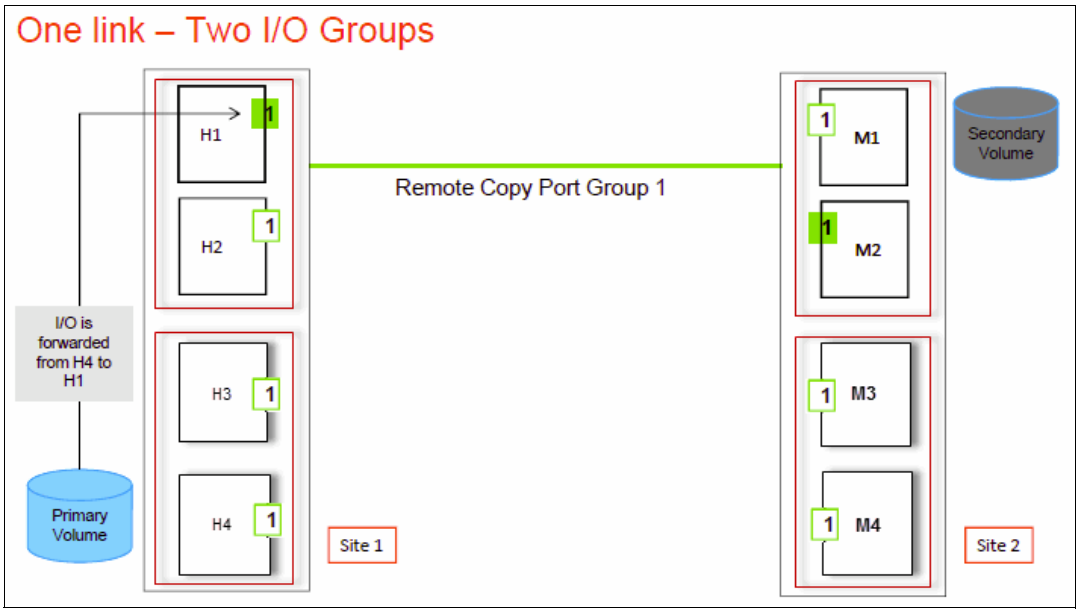

Figure 1-5 Dual I/O group with a single WAN link

1.3.3 Dual WAN links to a single I/O group

The design shown in Figure 1-6 shows a dual redundant WAN link connecting two Storwize V7000 clusters consisting of a single I/O group each. In this design, port group 1 and 2 have two IP addresses, each on a different node in a different system. This design enables two simultaneous IP replication sessions; however, there is only one active port per node. If any node should fail, the associated connection would also fail with the loss of half the bandwidth.

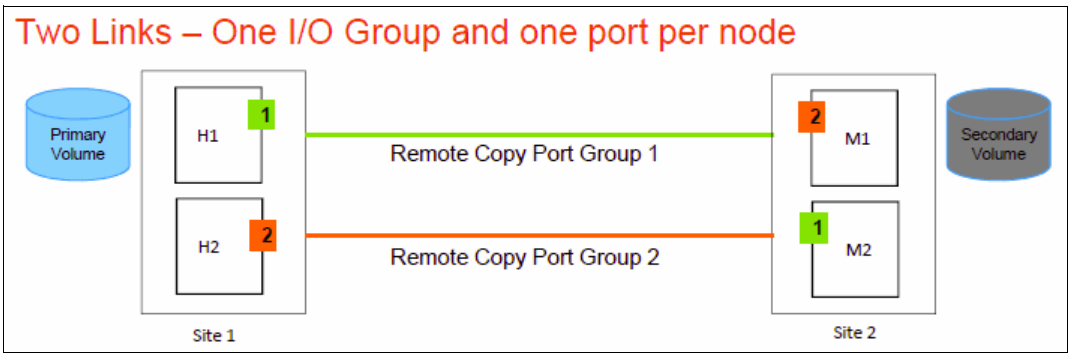

Figure 1-6 Dual WAN links to a single I/O group

1.3.4 Dual WAN links to dual I/O groups

If the required bandwidth means that there must be dual links, then a failure of one might lead to an effect on MM or GM if there is no redundancy at the node canister level. In this case, the advice would be to have two I/O groups at each site to enable recovery from any single node failure or outage, as shown in Figure 1-7.

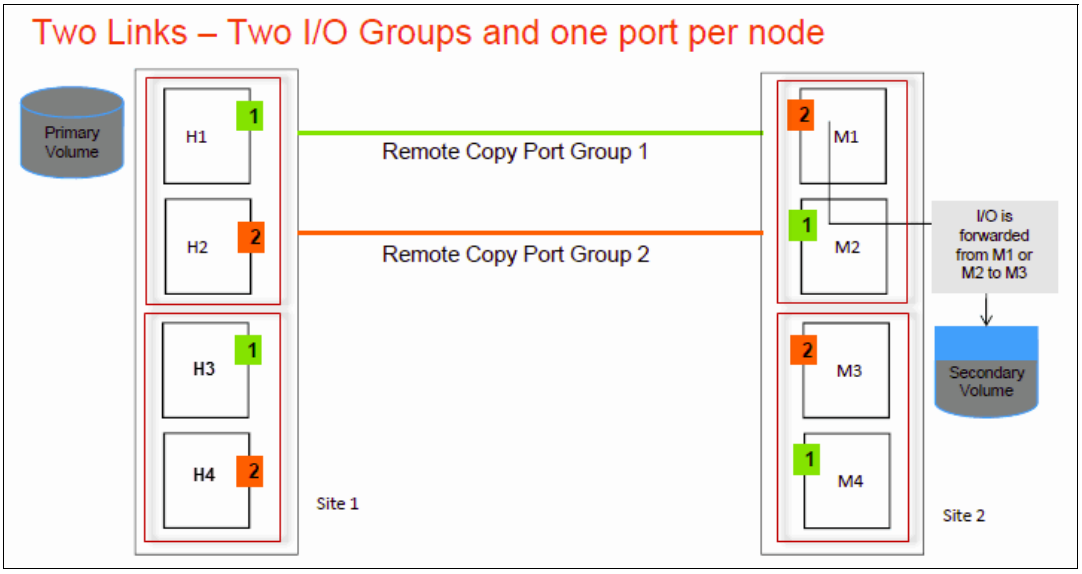

Figure 1-7 Dual WAN links to dual I/O groups

1.4 SVC/Storwize IP replication in more detail

The introduction of native IP replication provides an alternative to the traditional use of a pair of Fibre Channel over IP (FCIP) routers, or wave division multiplexors (WDMs) for WAN connectivity and to native Fibre Channel (FC) connectivity. Before the release of the version 7.2 code, if there was a need for replication, the only option was FC connectivity.

FC technology uses the concept of buffer-to-buffer (B2B) credits as a flow control method to maximize the performance of an individual link. The B2B credits are calculated depending on the speed and the distance of the link, and represent the number of frames a port can store. The faster the link, the more frames can be in flight, and the higher the B2B credit. Traditional IP does not have this concept.

1.4.1 Performance degradation with TCP/IP

Put simply, with TCP/IP, information transfer slows the farther you go. This is because of the latency caused by waiting for acknowledgment of each set of packets sent, because the next packet set cannot be sent until the previous one has been acknowledged:

```
Latency = Round Trip Time (RTT) for a single packet set
```

Figure 1-8 illustrates this packet flow for a typical IP link.

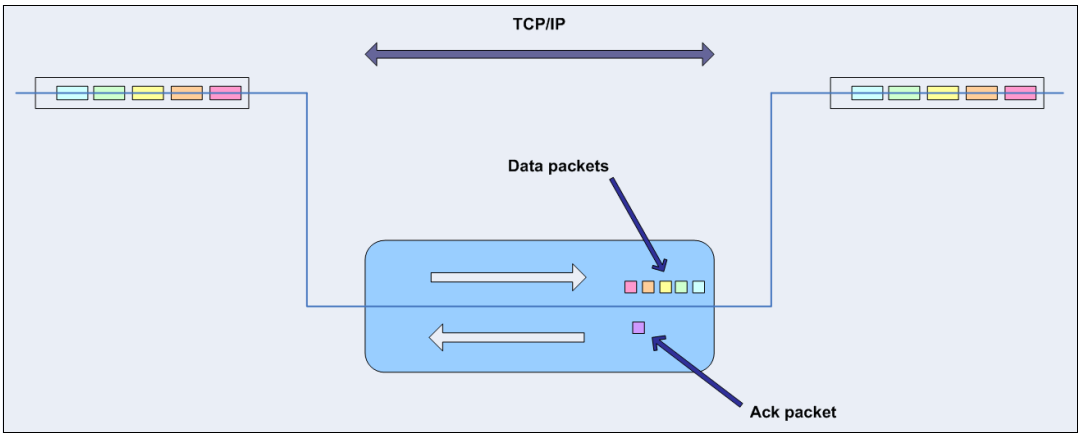

Figure 1-8 Link usage on a standard IP link

Figure 1-9 shows how throughput drops off as latency increases.

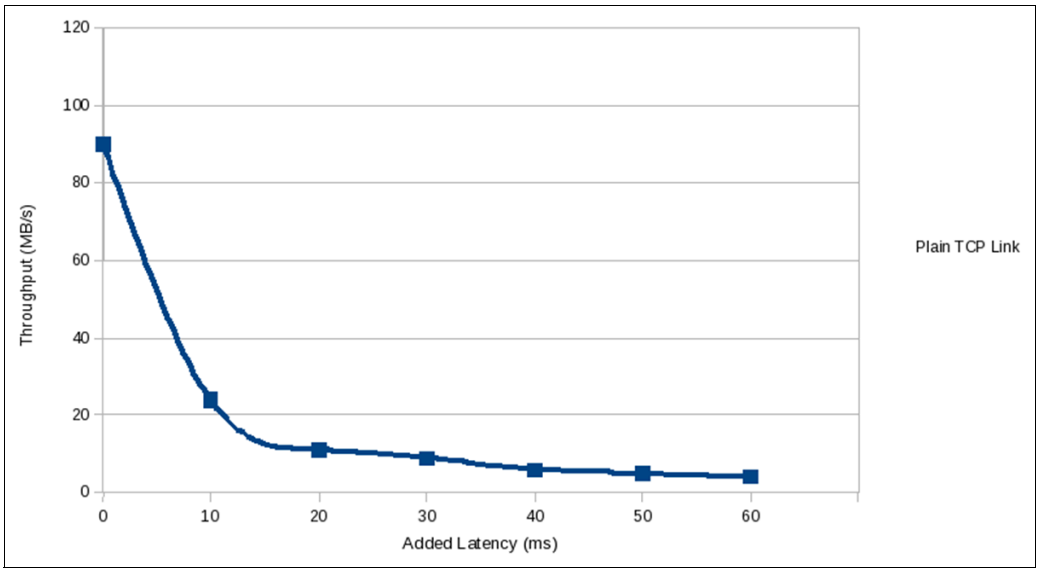

Figure 1-9 Effect of latency on performance

1.4.2 How SVC/Storwize improves performance

The technology built into the SVC/Storwize code uses TCP/IP latency to its advantage. Rather than wait for the acknowledgment to come back, it sends more sets of packets across other *virtual connections*. The number of virtual connections is controlled by the AI engine. This improves WAN connection use, which results in a data transfer rate approaching full line speed.

If packets are lost from any virtual connection, the data will be retransmitted, and the remote unit will wait for it. Presuming that this is not a frequent problem, overall performance is only marginally affected because of the delay of an extra round trip for the data that is resent.

Figure 1-10 illustrates data flow with multiple virtual connections.

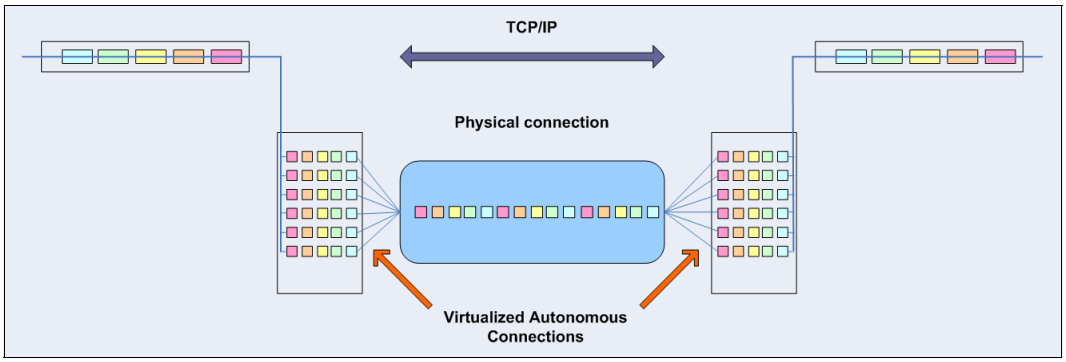

Figure 1-10 Link usage on an SVC/Storwize IP link

The AI monitors the link performance during data transfer, in addition to the memory and processor use of the node. It can adjust the number of virtual connections, the receive window size, and the packet size as appropriate to maintain optimum performance. This information is retained in the node so that, if the link is stopped and started again, it will restart with the previously learned settings.

What it does not do

SVC/Storwize IP replication does not manipulate the data in any way. It does not perform any of the following actions:

- Compress data
- Use deduplication
- Use User Datagram Protocol (UDP)
- Modify TCP/IP
- Use hard disk drive (HDD) caching

What it does do

It improves performance in several ways:

- Uses patented AI
- Fills the pipe
- Operates beneath the operating system
- Uses standard TCP/IP
- Is not affected by compressed and encrypted files

In summary, the SVC/Storwize solution is designed around the principle of using all of the pipe rather than changing the data. It works best with larger, more consistent amounts of data rather than bursty type data, which is another reason why we advocate GM/CV for IP replication. However, it also works with MM and GM.

1.4.3 Performance benefit

Figure 1-11 shows how the SVC/Storwize solution maintains near line-speed performance by masking the latency of the line. Even as the line latency increases, the performance of the technology enables the line to continue to exceed that of a plain link.

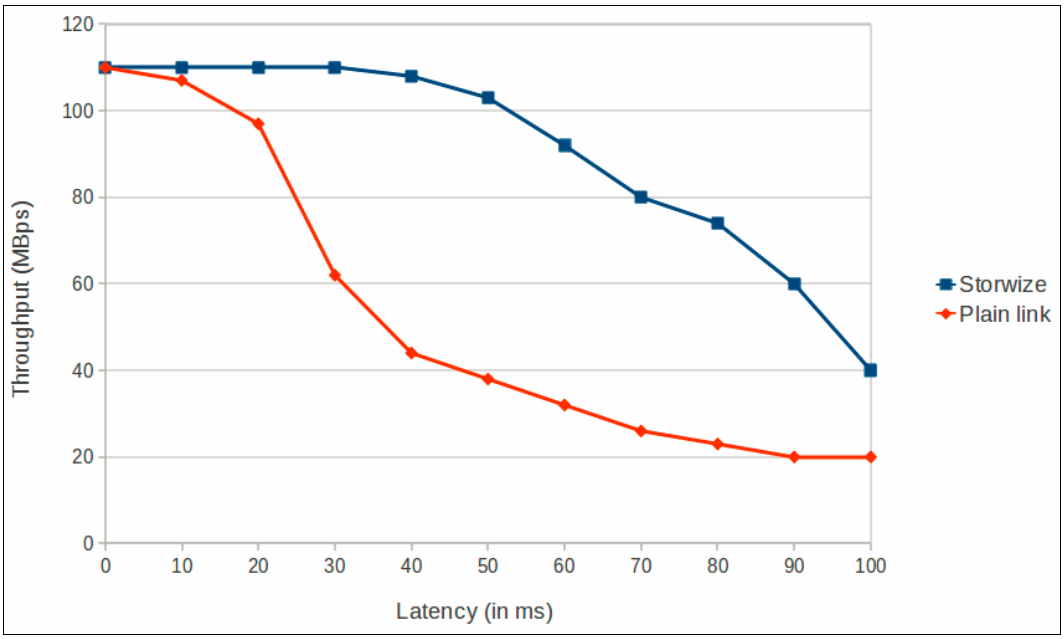

Figure 1-11 Storwize V7000 transfer rate comparison

2

IP replication configuration overview and considerations

This chapter describes the basic concepts and related steps for configuring Internet Protocol (IP) replication in IBM System Storage SAN Volume Controller (SVC) and IBM Storwize Family products. It is possible to set up IP replication between any IBM SVC/Storwize family member running code version 7.2 or later. It also documents the requirements to configure an IP partnership and intersite links.

2.1 IP partnerships

This topic provides information about items that relate to IP partnerships.

2.1.1 Intersite link planning

There are specific intersite link requirements that must be met when planning to use IP replication for remote copy.

If you use IP replication, you must meet the following requirements:

- Transmission Control Protocol (TCP) ports 3260 and 3265 are used by systems for IP partnership communications; therefore, these ports will need to be open. Port 3260 is used for link initialization and system discovery, and port 3265 is used for the IP replication traffic.
- The maximum supported Round Trip Time (RTT) between systems is 80 milliseconds (ms) for a link.
- For IP partnerships, the advised method of copying is Global Mirror with Change Volumes (GM/CV). This method is advised because of the performance benefits when used in combination with the SVC/Storwize IP replication technology (which works best with streaming data rather than bursty data). It is also advised because of the fact that it removes the local copy entirely from wide area network (WAN) glitches, and because it also smooths input/output (I/O) peaks and enables more efficient use of the intersite link.
- The amount of intersite heartbeat traffic is one megabit per second (Mbps) per link.
- The minimum bandwidth requirement for the intersite link is 10 Mbps. This, however, scales up with the amount of host I/O that you choose to do, as shown in Figure 2-1.

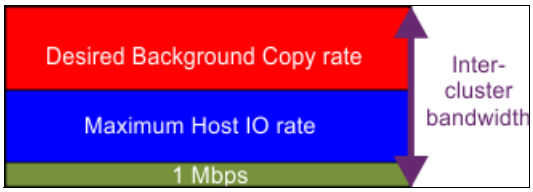

Figure 2-1 Intersite bandwidth requirements

The following equation describes the approximate minimum bandwidth required between two systems with less than (<) 5 ms round-trip time and an errorless link:

```
Minimum intersite link bandwidth in Mbps > required background copy in Mbps +
maximum host I/O in Mbps + 0.125 Mbps heartbeat traffic
```

If the link is dedicated and using only GM/CV traffic, then there will be no host I/O. It will all be sync I/O, and the link bandwidth formula changes.

Increasing latency and errors results in a higher requirement for minimum bandwidth. Network address translation (NAT) is not supported.

Important: It is important that both ports 3260 and 3265 are opened up through any firewalls for IP replication to be configured successfully. Additionally, ensure that there is no NAT in the intersite link.

2.1.2 IP partnership requirements

There are several requirements that must be remembered when creating IP partnerships:

- You can only create IP partnerships between two systems.
- A system can only be part of one IP partnership.
- You can create IP partnerships between any two SVC/Storwize products on the same layer. In other words, both systems must be at the storage layer or both systems must be at the replication layer. By default, all SVC systems are configured at the replication layer, and this *cannot* be changed. All Storwize systems by default are configured at the storage layer, but this *can* be changed.

  There are many other IBM Redbooks publications on the Storwize family that explain the concept of layers in more detail, for example, *Implementing the IBM Storwize V7000 V7.2*, SG24-7938 and *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574.
- You cannot use link-local addressing.
- If you use IPv4 addressing, the management IP addresses on both systems must be IPv4-compliant, and these addresses must have connectivity with each other.
- If you use IPv6 addressing, the management IP addresses on both systems must be IPv6-compliant, and these addresses must have connectivity with each other.
- If your system has IPv4 and IPv6 addressing on the ports, you must configure all remote copy links between the local and remote site for one of those types of addresses.
- You can configure ports from a maximum of two I/O groups from each system for an IP partnership.
- A node can have only one port configured in an IP partnership, either port 1 or 2, but not both. Where the optional 10 Gb Ethernet (GbE) card is installed in a system, ports 3 and 4 are also available.
- If your system has 10 Gbps and 1 Gbps Ethernet ports, you must configure all remote copy links between the local and remote site for just one of those speeds.
- Both ends of a partnership must be the same speed.
- A system can have simultaneous partnerships using Fibre Channel over IP (FCIP), but with separate systems.
- The IP partner systems must not be visible to each other over FC or Fibre Channel over Ethernet (FCoE) connections. FC/FCoE ports of one system should not be listed on the other when viewing the output of the **`lsfsabric`** command.
- Clusters configured in active IP partnerships should not be zoned with each other for FC/FCoE.
- Internet Small Computer System Interface (iSCSI) hosts *can* access volumes over IP ports participating in an IP partnership, but for the best performance it is advised to use separate ports for each.
- Virtual local area network (VLAN) tagging of the IP addresses configured for remote copy is supported starting with code level 7.4.0.x.
- If you configure two intersite links, you must configure separate remote-copy port groups, one for each link.
- If you have one intersite link, you can configure only one remote-copy port group for that link.
- No more than two intersite links and two remote-copy port groups are supported.

- If you have one remote-copy port group, configure one port from each node in one I/O group in that remote-copy port group.
- For systems with more than one I/O group, ports from a second I/O group can be added to the same remote-copy port group.
- If you have two remote-copy port groups and just a single I/O group, then on each system, configure one port from one node in the first remote-copy port group, and then a port from the other node in the second remote-copy port group.
- For systems with more than one I/O group, ports from a second I/O group can be added to each of the two remote-copy port groups.

  For more details about I/O groups, port groups, and node port configurations, see 1.3, “Example scenarios” on page 4.
- If you connect systems by directly attaching them without switches, you must have only two direct-attach links, and both direct-attach links must be on the same I/O group. You should use two port groups, where a port group should only contain the two ports that are directly linked.

2.1.3 Configuring IP partnerships

The procedure to configure a partnership with a remote system using IP replication is specified here.

To configure a partnership with a remote system by using IP replication, follow these steps:

1. Determine whether you have met all partnership requirements, as described in 2.1.2, “IP partnership requirements” on page 13.
2. For each node port that is to be used in the IP partnership, configure the ports on both systems using the command-line interface (CLI) **`cfgportip`** command. Configuration includes assigning each port to a remote-copy port group, as described in section 2.1.4, “Remote-copy port groups” on page 15. Remote-copy port groups are unique to IP partnerships, and are the local and remote IP addresses accessible to each other through an IP partnership.

   You can also use the management graphical user interface (GUI) to do this task:

   a. Select **Settings** → **Network** → **Ethernet Ports**.
   b. Right-click the port, and then select **Modify**.
3. Configure IP replication on the local system by running the **`mkippartnership`** command on that system. Specify the Challenge Handshake Authentication Protocol (CHAP) secret, if necessary. See section 2.1.7, “Using CHAP with IP partnerships” on page 16 for more information about CHAP. The partnership link bandwidth parameter must be less than or equal to (<=) the bandwidth of the intersite link.
4. If you want to do background copy, the amount of link used is governed by the background copy rate parameter. Adjust this parameter so that there is enough bandwidth available for the largest host I/O burst expected. Use the following equation to determine the approximate percentage:

   ```
   background copy rate = ((intersite links bandwidth in MBps - approximate
   maximum host I/O in MBps expected - 0.125 MBps heartbeat traffic) / Intersite
   links bandwidth in MBps) x 100
   ```

   You can also use the management GUI to do this task by selecting **Copy Services** → **Partnerships** → **Create Partnership**.
5. Configure IP replication on the remote system by running the **`mkippartnership`** command on the remote system, specifying the CHAP secret that you provided previously, if any.

For a detailed step-by-step guide to configuring IP partnerships, see Chapter 3, “Implementing IP replication in the SVC/Storwize family” on page 21.

2.1.4 Remote-copy port groups

Remote-copy port groups are unique to IP partnerships. They are the local and remote IP addresses accessible to each other through an IP partnership. A remote-copy port group must contain at least one IP address in the local system and one IP address in the remote system. You must configure individual ports on the local and remote systems and assign to them a remote-copy port group number before establishing the IP partnership.

A remote-copy port group is designated by a number 0 - 2 that identifies a set of IP addresses. These port groups are defined as follows:

Group 0 These ports are not configured for remote copy (supports iSCSI host attachment by default).

Group 1 Ports that belong to remote-copy port group 1.

Group 2 Ports that belong to remote-copy port group 2.

The system will select one pair of IP addresses within port groups 1 and 2, one address on each system, and will open a remote copy connection between them. It is not possible to choose which IP addresses will be selected, it is a system choice based on the IP address of the port.

See 3.1.3, “Configure the Ethernet ports on the partner systems” on page 24 for a walk-through describing how to assign IP addresses to ports, and how to assign those ports to port groups.

Tip: Each IP address can be shared for iSCSI host attachment and remote copy functionality. Therefore, the appropriate settings must be applied to each IP address. For optimal performance, consider isolating iSCSI and IP replication traffic on separate ports.

2.1.5 Node port status

After an IP partnership has been defined, it is possible to determine which ports are in use in a port group and which are not. Using the **`lsportip`** command will show the status of all ports. In the output, the value of the **`remote_copy_status`** field for IPv4 ports, or the **`remote_copy_status6`** field for IPv6 ports, will be `used`. This is true only if the port has been configured for use in the IP partnership, and has been selected by the system for sending and receiving remote copy data.

If a port has been configured for use in the IP partnership, but is not currently selected for sending and receiving remote copy data, the value of the corresponding field will be `unused`. If a port has not been configured for use by the IP partnership, this field will be blank.

Example 2-1 on page 16 shows the output from one of the SVCs used for this paper. There is one I/O group, and two ports (port 2 on each node) are configured for use in IP replication.

There are no unused ports in this example, so there are two intersite links in use, and therefore two port groups. All other ports are not part of any I/O group, because they are blank.

Example 2-1 shows the **lsportip** command output.

Example 2-1 The lsportip command output

```
IBM_2145:ITSO_Local_Cluster:superuser>lsportip
id node_id node_name IP_address  mask     gateway   IP_address_6 prefix_6 gateway_6 MAC        duplex
state        speed failover link_state host remote_copy host_6 remote_copy_6 remote_copy_status remote_copy_status_6
1  1     node1                                                    e4:1f:13:30:19:e4 Full
unconfigured 1Gb/s no       active          0                  0
1  1     node1                                                    e4:1f:13:30:19:e4 Full
unconfigured 1Gb/s yes      active          0                  0
2  1     node1   192.168.1.10 255.255.255.0 192.168.1.1                 e4:1f:13:30:19:e6 Full
configured   1Gb/s no       active     no   1                  0            used
2  1     node1                                                    e4:1f:13:30:19:e6 Full
configured   1Gb/s yes      active          0                  0
1  2     node2                                                    e4:1f:13:b7:7f:78 Full
unconfigured 1Gb/s no       active          0                  0
1  2     node2                                                    e4:1f:13:b7:7f:78 Full
unconfigured 1Gb/s yes      active          0                  0
2  2     node2   192.168.1.11 255.255.255.0 192.168.1.1                 e4:1f:13:b7:7f:7a Full
configured   1Gb/s no       active     no   2                  0            used
2  2     node2                                                    e4:1f:13:b7:7f:7a Full
configured   1Gb/s yes      active          0                  0
```

2.1.6 IP replication link failover

As previously mentioned in 1.3, “Example scenarios” on page 4, only one port from a port group on each partner system can be active at any one time. That is why in Example 2-1, you can deduce that there are in fact two replication links. Because there are two ports in the `used` state, it means that there must be two port groups and therefore two intersite links. In this scenario, if one link were to fail for some reason, either through the loss of the WAN link, a node fault, or a port fault, the entire link would drop and be lost.

Communication would continue through the remaining link. The failed link would not re-establish until the fault that caused link loss was rectified, because there are no other ports defined in that port group for failover. If there were a second I/O group, then the port group definition would follow that, as outlined in 1.3.4, “Dual WAN links to dual I/O groups” on page 7.

In this case, a loss of a node would lead to a temporary loss of the link, while the discovery process is triggered and re-establishes the link using the failover node. Typically this process takes around 30 - 60 seconds to complete. Clearly, with only one intersite link, any failure, even with failover node ports defined in the port group, would result in a temporary outage to the intersite connection.

With Metro Mirror (MM) and GM volumes, this might result in the relationships stopping, and the need for manual intervention to restart them. This is one of the reasons why GM/CV is the preferred option for IP replication, because it can be more tolerant of short outages such as this if only one intersite link is available.

2.1.7 Using CHAP with IP partnerships

Data exchange between the local system and partner system over an IP connection can be protected through CHAP, which uses a shared secret to authenticate systems with each other when sending requests.

Note that you can also specify that the same CHAP secret be used to authenticate the system with iSCSI-attached hosts. The system-wide CHAP secret is used for all CHAP authentication from the local system to partner systems, and to iSCSI-attached hosts. For more information about CHAP, see Appendix B, “CHAP Information and Implementation” on page 57.

2.2 Setting up remote copy relationships on SVC/Storwize

Establishing remote copy relationships on a Storwize or SVC system is already covered in several existing publications that are available from IBM.

For example, see *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574, available on the following web page:

`http://www.redbooks.ibm.com/redbooks/pdfs/sg247574.pdf`

Figure 2-2 represents a high-level conceptual view of how GM/CV isolates host volumes from the replication process.

Figure 2-2 Global Mirror with Change Volumes concept

For details about how to set up IP replication using GM/CV, see Chapter 3, “Implementing IP replication in the SVC/Storwize family” on page 21.

2.3 Migrating between FC and IP replication

It is possible to move existing mirror relationships between different replication partnerships. In each case this requires an outage to stop host I/O to the volumes in the relationships to be converted from one type of infrastructure to the other.

The procedures outlined below preserve the data on the volumes and removes the need to do a complete re-sync on the mirror volumes when the new relationships are established, you may decide that it is simpler to delete the existing relationships, create new ones and perform a complete re-sync from scratch. The following conversions are available:

- FC to IP and vice versa
- 1 Gb links to 10 Gb links and vice versa
- IPv4 to IPv6 and vice versa

2.3.1 Migration between FC and IP, or IP and FC

When migrating from or to FC partnerships, to ensure that all cached data on the host is flushed to the volumes, make sure that any file systems that are mounted on replicated volumes are unmounted first. This process must be done for all relationships. It must also be done on all hosts if the same volume is used by multiple hosts, for example using clustered file systems, such as VMware Virtual Machine File System (VMFS).

If you are using other applications, you must ensure that you synchronize all cached data from application to disk. This procedure might be application-specific as well. For example, Oracle, IBM DB2®, and so on, might need to be stopped, although for some applications you might need to run synchronization on the host.

To prepare for the migration, follow these steps:

1. If there are GM/CV remote-copy relationships, complete the following steps. If not, go to the next step:
   a. Stop the relationship and change to non-cycling GM.
   b. Start the relationship.

   Important: Make sure that the relationship makes the transition to `consistent_synchronized` (or else wait until the relationship's state is `consistent_synchronized`).

2. Stop the relationships without the **`-access`** flag and verify that the status for each relationship is `in_sync`.
3. Delete the remote-copy relationships.
4. After all remote-copy relationships are deleted, stop and delete the partnerships from both systems.

FC to IP migration

For FC to IP migration, follow these additional steps:

1. Remove/delete the zoning between the two sites so that the two systems are not listed as available systems. Remember when using IP replication that neither system should be able to see the other before the creation of the IP replication partnership.
2. Configure the IP ports and establish IP partnerships. Additionally, configure CHAP if required.
3. Create the remote-copy relationships with the **`-sync`** flag for MM, GM, or GM/CV with the original master and auxiliary volumes (previously used in the FC partnerships) on respective sites.
4. Add the change volumes to the respective relationships.
5. Start the remote-copy relationship(s).

IP to FC migration

For IP to FC migration, follow these additional steps:

1. Unconfigure the ports in the remote copy port groups (set them to 0) on both systems.
2. Create the zones between the two sites so that the two systems are listed as available systems when you run **`lspartnershipcandidate`** and **`lsfabric`**.
3. Create the remote-copy relationships with the **`-sync`** flag for MM, GM, or GM/CV with the original master and auxiliary volumes (previously used in the FC partnerships) on respective sites.
4. Add the change volumes to the respective relationships.
5. Start the remote-copy relationship(s).

These steps complete the migration of remote-copy relationships from FC to native IP or vice versa.

2.3.2 Migration between IPv4 and IPv6, or IPv6 and IPv4

To migrate from IPv4 to IPv6, the following requirements must be met:

- System IPs have IPv6 addresses configured.
- Data path IPs (IPs configured by using **cfgportip**) have IPv6 addresses.

> **Restriction:** You can *assign* IPv6 addresses to ports while IP partnerships are active. However, you cannot *add* them to remote-copy port groups.

Complete the following procedure:

1. Stop the relationships without the **-access** flag and verify that the status for each relationship is `in_sync`.
2. After all remote-copy relationships are stopped, stop the IP partnership on both systems.
3. Add the IPv6 IP addresses, which are configured on data path IP Ports (by using the **cfgportip** command), in respective remote copy port groups.

> **Remember:** This step causes your remote-copy status to change from `used` to `unused`, because discovery for the new paths on new IP addresses has not yet happened.

4. Modify the IP partnerships to do discovery over IPv6 addresses.

> **Confirm:** Because new data paths over IPv6 addresses are now available, partnership will first change to `not_present`, and then to `fully_configured`. If it remains in `not_present`, monitor the Event Log to see if an error is getting triggered, and run the appropriate Directed Maintenance Procedure (DMP).

5. Start the remote-copy relationships.

This procedure completes the migration of remote-copy relationships from IPv4 to IPv6. This same procedure can be applied when migrating from IPv6 to IPv4 by applying suitable substitutes for IPv4 rather than IPv6.

2.3.3 Migration between 1 Gbps and 10 Gbps, or 10 Gbps and 1 Gbps

Before you attempt to migrate existing deployments from 1 Gbps links to 10 Gbps links, see the limitations and considerations in IP partnership configuration, 2.1.2, “IP partnership requirements” on page 13.

> **Restriction:** You cannot mix link speeds. For example, if there are two links, both links should either be 10 Gbps links or both should be 1 Gbps links.

Complete the following procedure:

1. Stop the relationships without the **-access** flag and verify that the status for each relationship is *in_sync*.
2. After all remote-copy relationships are stopped, stop the IP partnership on both systems.
3. Add the IP addresses, which are configured on data path IP ports on 10 Gbps links in respective remote-copy port groups.

4. Remove the existing 1 Gbps ports from the remote-copy groups.

   Remember: This step causes your remote-copy status to change from `used` to `unused`, because discovery for the new paths on new IP addresses has not yet happened.

5. Start the IP partnerships.

   Confirm: Because new data paths over IPv6 addresses are now available, partnership will first change to `not_present`, and then to `fully_configured`. If it remains in `not_present`, monitor the node error/DMP if it is getting triggered, and examine the appropriate DMP.

6. Start the remote-copy relationships.

This procedure completes the migration from 1 Gbps links to 10 Gbps links. This same procedure can be applied when migrating from 10 Gbps to 1 Gbps by applying suitable substitutes for 1 Gbps ports rather than 10 Gbps ports.

3

Implementing IP replication in the SVC/Storwize family

This chapter covers the steps required to configure Internet Protocol (IP) replication in the IBM System Storage SAN Volume Controller (SVC)/IBM Storwize family of products using the web-based graphical user interface (GUI). A detailed description of the example environment is provided. For details about using the command-line interface (CLI) to configure IP replication, see Appendix A, “Command-line interface” on page 53.

The following sections will describe the steps to configure an IP replication partnership between two SVC clusters, and to create and initiate a remote-copy relationship of volumes between these partner systems.

3.1 Configuring IP replication

The following sections describe how to configure IP replication.

3.1.1 Overview of required steps

Before configuring IP replication, ensure that the following prerequisites have been met:

1. Confirm the presence and details of the intersite IP links to be used for replication traffic between your partner systems. For more details about the requirements for intersite links, see Chapter 2, “IP replication configuration overview and considerations” on page 11.
2. Determine the number of Ethernet ports and replication port groups to be used.
3. Obtain IP addresses for the IP ports to be used, and ensure that the intersite links support the routing of traffic between the addresses on each end, and that Transmission Control Protocol (TCP) ports 3260 and 3265 are open for traffic between the partner systems.

To configure IP replication, the following steps must be performed:

1. Configure the Ethernet ports on each of the partner systems.
2. Create the IP partnership on each of the partner systems.
3. Create the wanted remote-copy relationship(s).

3.1.2 Physical details of an example environment

The configuration examples reference a lab environment containing two SVC clusters. Each SVC cluster contains two nodes. Our intersite connectivity consists of two 1 gigabits per second (Gbps) links. For the IP replication traffic, the example uses port 2 on each SVC node.

The example configuration in Figure 3-1 shows the two SVC clusters used to demonstrate IP replication in this paper.

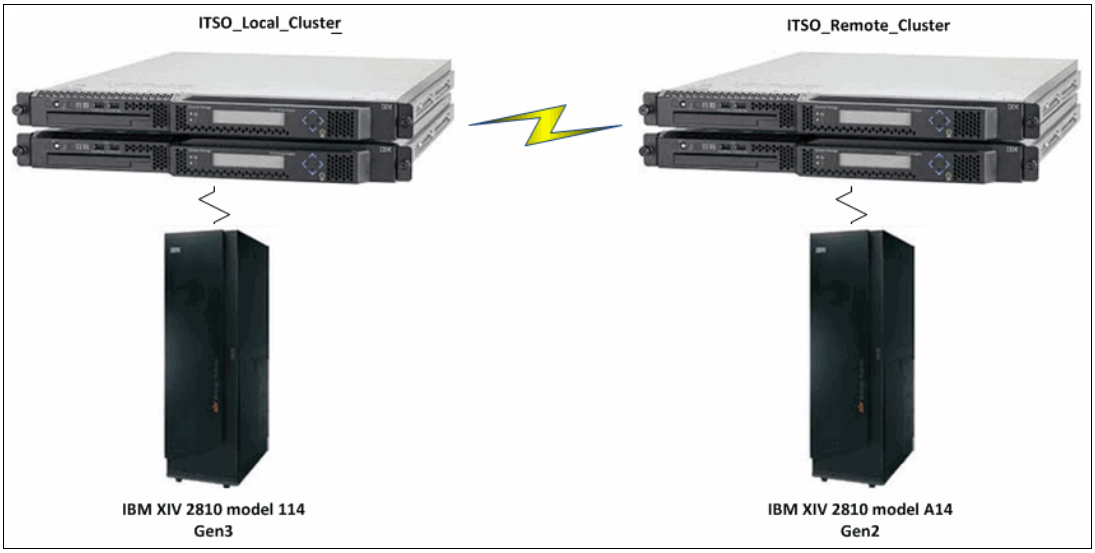

Figure 3-1 Example configuration

This example highlights the following elements of the configuration:

- One SVC named ITSO_Local_Cluster and one SVC named ITSO_Remote_Cluster. Both systems use IBM XIV Storage Systems for back-end storage.
- This example uses network port 2 on each of the SVC nodes for IP replication, with two 1 Gbps intersite links.
- Because Global Mirror with Change Volumes (GM/CV) is the advised method to use with IP replication, GM/CV is used in the following examples.

The following sections provides details about the SVC clusters.

ITSO_Local_Cluster

ITSO_Local_Cluster has the following characteristics:

Cluster IP address 9.32.248.174

Node 1 port 2

IP address	192.168.1.10
Subnet mask	255.255.255.0
Gateway	192.168.1.1
Remote port group	Group 1

Node 2 port 2

IP address	192.168.1.11
Subnet mask	255.255.255.0
Gateway	192.168.1.1
Remote port group	Group 2

Master volumes (17 gigabytes, or GB)

repl_master_1
repl_master_2
repl_master_3
repl_master_4

Master change volumes (17 GB)

repl_master_cv_1
repl_master_cv_2
repl_master_cv_3
repl_master_cv_4

ITSO_Remote_Cluster

ITSO_Remote_Cluster has the following characteristics:

Cluster IP address 9.32.248.89

Node 1 port 2

IP address	192.168.1.12
Subnet mask	255.255.255.0
Gateway	192.168.1.1
Remote port group	Group 1

Node 2 port 2

IP address	192.168.1.13
Subnet mask	255.255.255.0
Gateway	192.168.1.1
Remote port group	Group 2

Auxiliary volumes (17 GB)

repl_auxiliary_1
repl_auxiliary_2
repl_auxiliary_3
repl_auxiliary_4

Auxiliary change volumes (17 GB)

repl_auxiliary_cv_1
repl_auxiliary_cv_2
repl_auxiliary_cv_3
repl_auxiliary_cv_4

Information: The IP addresses (IPs) used in this scenario are just one example of what IPs can be used. Any range of IPs are allowable, private networks included, if they conform with your own IP network and the requirements described in 2.1.2, “IP partnership requirements” on page 13.

3.1.3 Configure the Ethernet ports on the partner systems

This section describes the steps to configure the IP ports on the partner systems.

Remember: In the lab environment used for the examples in this section, Ethernet port 2 is configured on each of the two SVC nodes on the partner systems. Alternative Ethernet port configurations are supported. See 2.1, “IP partnerships” on page 12 for more information about the various supported configurations.

Configuring the Ethernet ports on ITSO_Local_Cluster

To configure the IP ports on ITSO_Local_Cluster, follow these steps:

1. Access the **Network** pane in the web-based GUI using the **Settings** icon in the left navigation menu, as show in Figure 3-2.

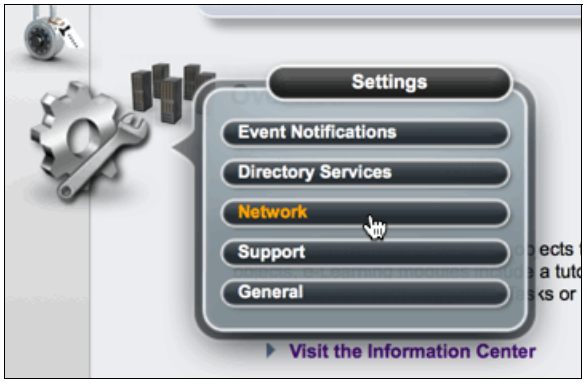

Figure 3-2 Accessing the Network pane in the web-based GUI

2. On the Network pane, click **Ethernet Ports** in the left Network menu, as shown in Figure 3-3.

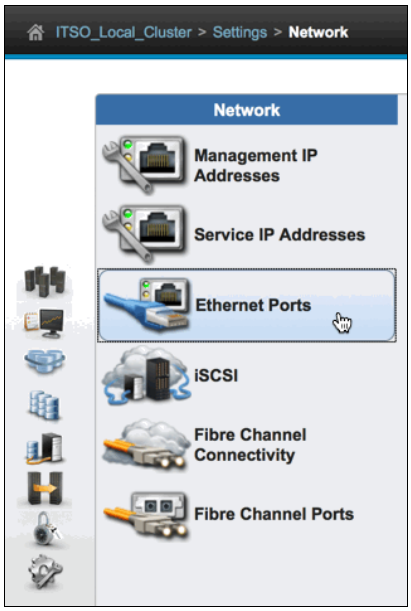

Figure 3-3 Accessing the Ethernet Ports pane in the web-based GUI

3. Right-click the row of the wanted port and select **Modify** to modify its IP settings, as shown in Figure 3-4. In this example, configure port 2 on node 1.

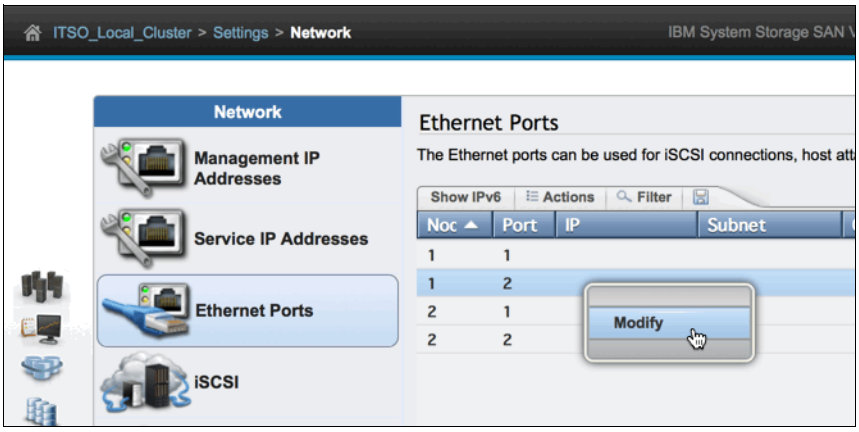

Figure 3-4 Modifying Ethernet Port settings in the web-based GUI

4. In the Modify Port dialog window shown in Figure 3-5 on page 26, enter the appropriate values for **IP address**, **Subnet mask**, and **Gateway**. In this example, the following values were used for port 2 on node 1:

IP address	192.168.1.10
Subnet mask	255.255.255.0
Gateway	192.168.1.1
iSCSI hosts	Disabled
Remote copy	Group 1

5. Additionally, select or clear the **iSCSI hosts** box to enable or disable Internet Small Computer System Interface (iSCSI) host traffic on this port. In this example, clear the box to disable iSCSI traffic on this port.

Tip: Although the sharing of IP ports for both IP replication traffic and iSCSI host traffic is supported, for performance purposes, it is advised to isolate IP replication traffic. For more information, see Chapter 2, “IP replication configuration overview and considerations” on page 11.

6. To enable the port for IP replication, use the **Remote copy** radio buttons to select a port group. In this example, choose **Group 1** for Ethernet port 2, as shown in Figure 3-5.

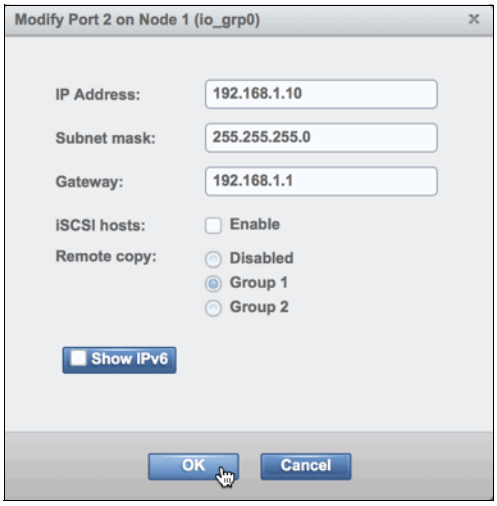

Figure 3-5 Configuring Ethernet Port settings in the web-based GUI

7. When you click **OK**, the Modify Ethernet Port IP address dialog window will display the task status for modifying the Ethernet port IP address settings, as shown in Figure 3-6. After the task completes, click **Close** to close the task status window.

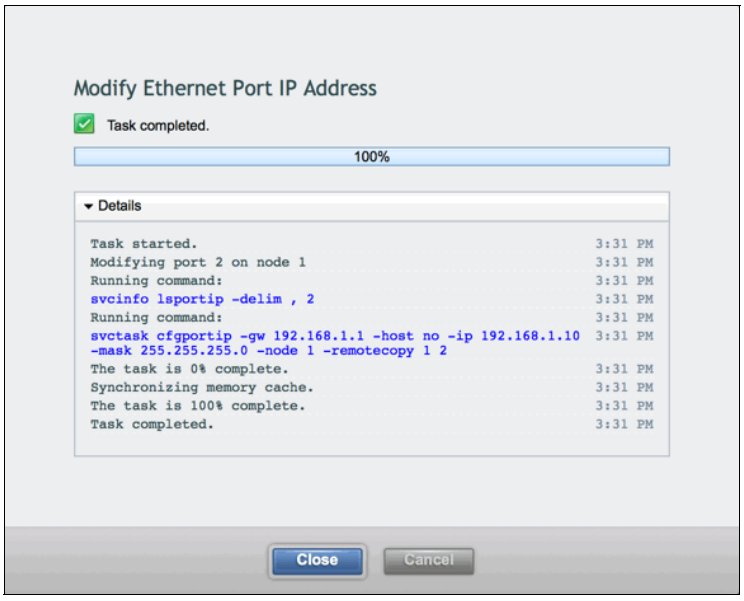

Figure 3-6 Task status for Ethernet port IP address modification in the web-based GUI

8. The **Ethernet Ports** pane should now reflect the changes made, as shown in Figure 3-7. Ensure that your settings are correct.

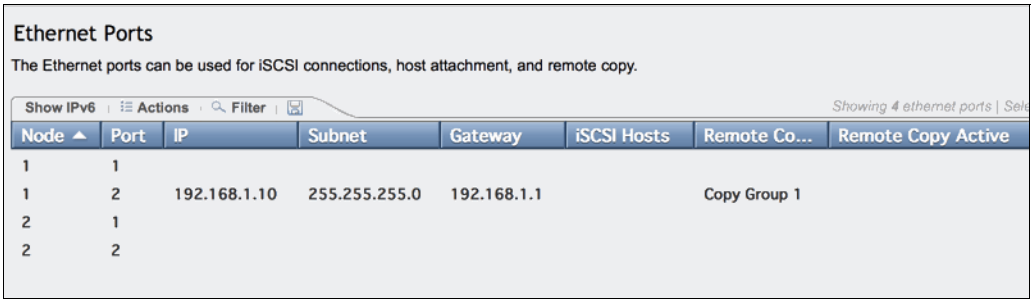

Figure 3-7 Verifying Ethernet Port settings for port 2 on node 1 in the web-based GUI

9. Repeat steps 3 on page 25 to 8 to configure the IP settings for port 2 on node 2. In this example, the following values were used for port 2 on node 2:

IP address	192.168.1.11
Subnet mask	255.255.255.0
Gateway	192.168.1.1
iSCSI hosts	Disabled
Remote copy	Group 2

10. The **Ethernet Ports** pane should now reflect the changes made to port 2 on node 2, as shown in Figure 3-8.

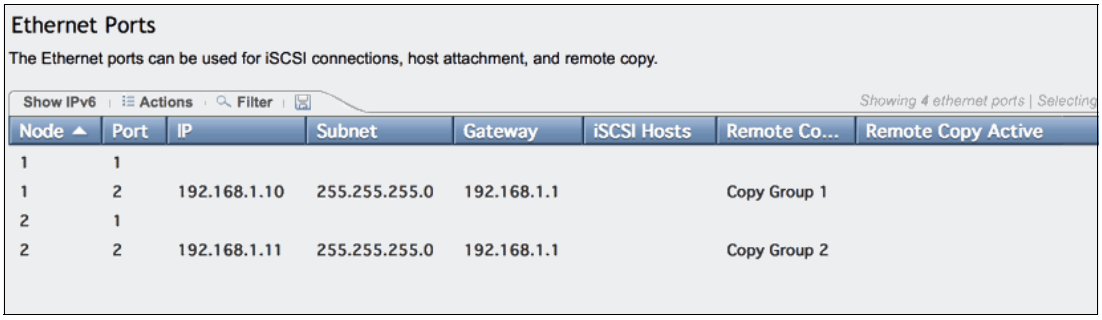

Figure 3-8 Verifying Ethernet Port settings for port 2 on node 2 in the web-based GUI

Configuring the Ethernet ports on ITSO_Remote_Cluster

To configure the IP ports on ITSO_Remote_Cluster, access the web-based GUI of ITSO_Remote_Cluster and repeat the procedure outlined in “Configuring the Ethernet ports on ITSO_Local_Cluster” on page 24, using the appropriate values. In this example, configure the Ethernet Ports on ITSO_Remote_Cluster using the following values:

Port 2 on Node 1

IP address	192.168.1.12
Subnet mask	255.255.255.0
Gateway	192.168.1.1
iSCSI hosts	Disabled
Remote copy	Group 1

Port 2 on Node 2

IP address	192.168.1.13
Subnet mask	255.255.255.0
Gateway	192.168.1.1
iSCSI hosts	Disabled
Remote copy	Group 2

Figure 3-9 shows the port configuration for the Ethernet Ports on ITSO_Remote_Cluster.

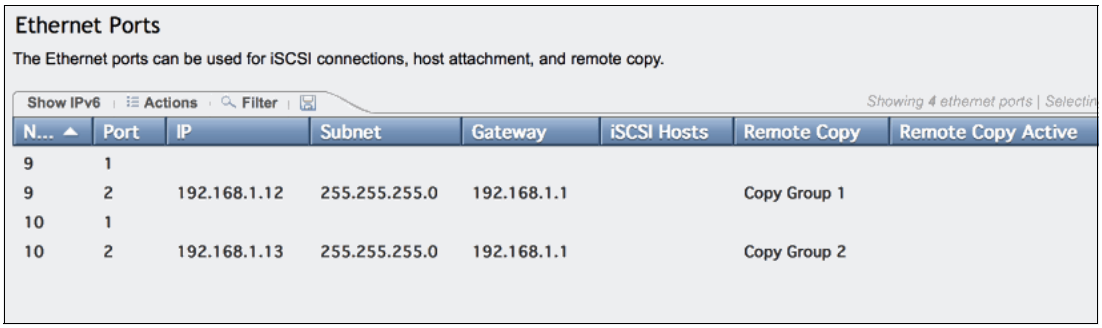

Figure 3-9 Verifying Ethernet Port settings for ITSO_Remote_Cluster in the web-based GUI

The Ethernet port settings are now complete, and you can proceed to configure the IP partnership.

3.1.4 Create and configure the IP partnership on each of the partner systems

In this section, we detail the steps to configure the IP partnership on the partner systems.

Remember: To create a Partnership between SVC and Storwize family systems, those systems must be in the same layer. SVC systems are always in the replication layer. Storwize systems are in the storage layer by default, but can be configured in the replication layer. To create a partnership between an SVC and a Storwize system, you must ensure that the Storwize system layer is set to replication. For more information, see 2.1.2, “IP partnership requirements” on page 13.

For information about troubleshooting IP partnership layer issues, see 4.2.2, “Error code CMMVC8354E” on page 48.

Configuring the IP partnership on ITSO_Local_Cluster

To configure the IP partnership on ITSO_Local_Cluster, perform the following steps:

1. Access the Partnerships pane in the web-based GUI using the **Copy Services** icon in the left navigation menu, as shown in Figure 3-10 on page 29.

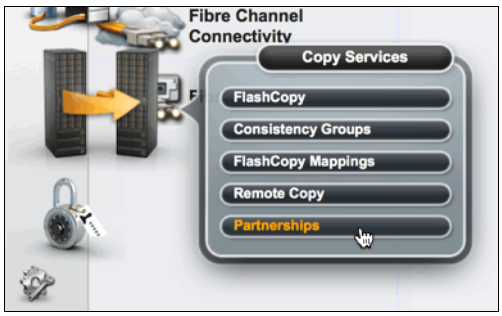

Figure 3-10 Accessing the Partnerships pane in the web-based GUI

2. On the Partnerships pane, click **Create Partnership**, as shown in Figure 3-11.

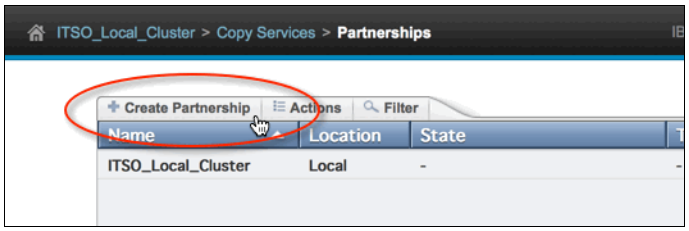

Figure 3-11 Creating a partnership in the web-based GUI

3. In the Create Partnership dialog window shown in Figure 3-12, select **IP**.

Figure 3-12 Selecting IP partnership replication type in the web-based GUI

4. In the Create Partnership dialog window shown in Figure 3-13 on page 30, enter values for **Partner system IP address**, **Link bandwidth**, **Background copy rate**, and optionally, **Partner system's CHAP secret**. In this example, the following values were used:

Partner system IP address	9.32.248.89
Link bandwidth in megabits per second (Mbps)	2000 (our example environment contains two 1 Gbps intersite links)
Background copy rate (%)	50 (we chose to use up to 50% of the link bandwidth for the initial background synchronization)
Partner system's CHAP secret	None (the example environment does not use CHAP authentication)

Tip: For the partner system IP address, enter the *management IP address* ("cluster IP") for the partner system. Do not enter the IP addresses that you configured for IP replication on the Ethernet ports of the partner system.

Figure 3-13 shows the example values.

Figure 3-13 Configuring IP partnership settings in the web-based GUI

5. When you click **OK**, the Create IP Partnership dialog window will display the task status for creating the IP partnership, as shown in Figure 3-14. When the task completes, click **Close** to close the task status window.

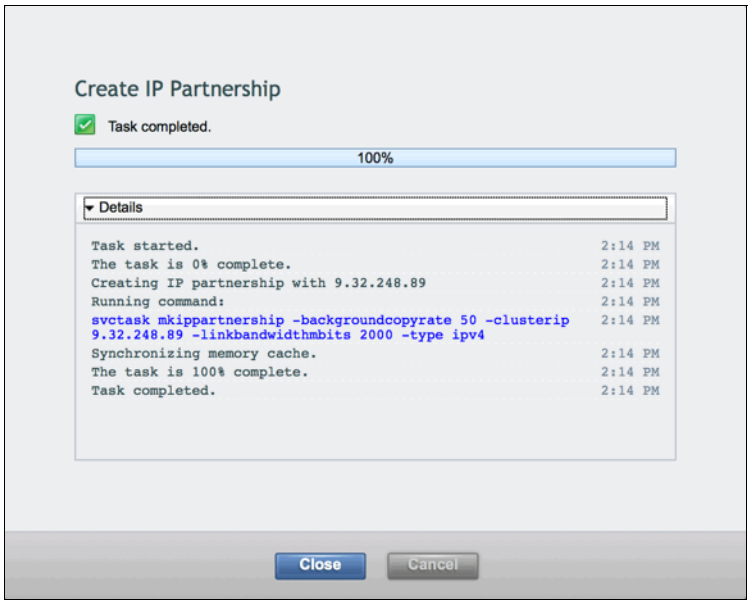

Figure 3-14 Task status for IP partnership creation in the web-based GUI

6. Confirm that the Partnerships pane displays the new partnership that you just created, as shown in Figure 3-15.

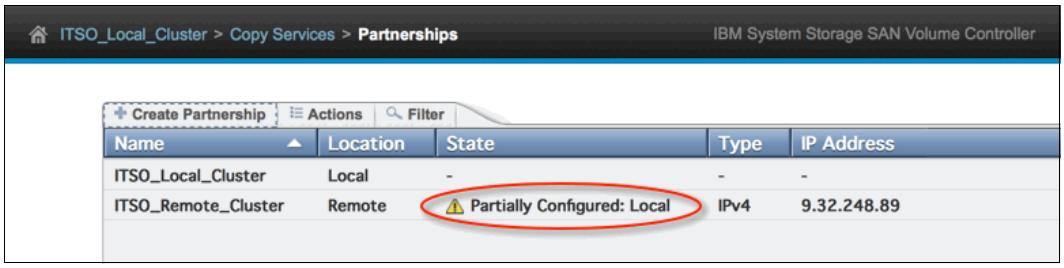

Figure 3-15 Verifying IP partnership creation in the web-based GUI

Note: Make note of the **State** of the partnership. When the partnership is successfully created on the first system, the **State** value on that system will display `Partially Configured: Local`. When the partnership is successfully created on the second system, the **State** on that system might initially display `Not Present`. After a brief period of time, the **State** value will display `Fully Configured` on each system if the IP partnership was successfully created on each partner system.

For further information about troubleshooting IP partnership creation, see Chapter 4, “Best practices and troubleshooting” on page 45.

Configuring the IP partnership on ITSO_Remote_Cluster

To configure the IP partnership on ITSO_Remote_Cluster, access the web-based GUI of ITSO_Remote_Cluster and repeat the procedure outlined in “Configuring the IP partnership on ITSO_Local_Cluster” on page 28, using the appropriate values. In this example, configure the IP partnership on ITSO_Remote_Cluster using the following values:

Partner system IP address 9.32.248.174

Link bandwidth (Mbps) 2000 (the example environment contains two 1 Gbps intersite links)

Background copy rate (%) 50 (we chose to use up to 50% of the link bandwidth for the initial background synchronization)

Partner system’s CHAP secret None (our example environment does not use CHAP authentication

After initially configuring the IP partnership on ITSO_Remote_Cluster, the **Partnerships** pane shows the state of this partnership as Not Present, as shown in Figure 3-16.

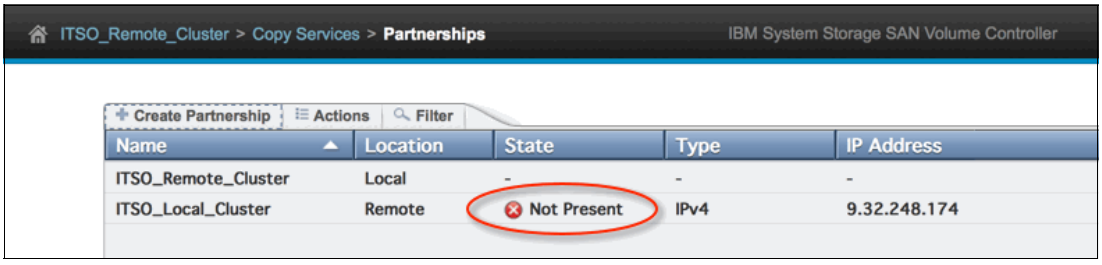

Figure 3-16 Initial IP partnership state on ITSO_Remote_Cluster in the web-based GUI

After a few moments, the State of the IP partnership should display *Fully Configured* on both systems, as shown in Figure 3-17 and Figure 3-18.

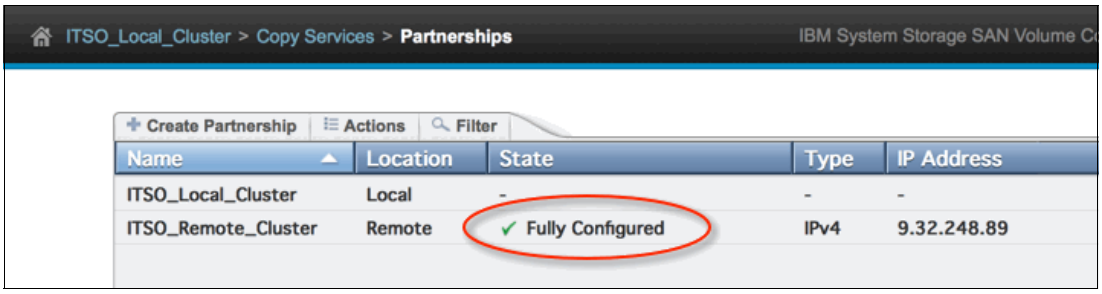

Figure 3-17 Fully configured IP partnership on ITSO_Local_Cluster in the web-based GUI

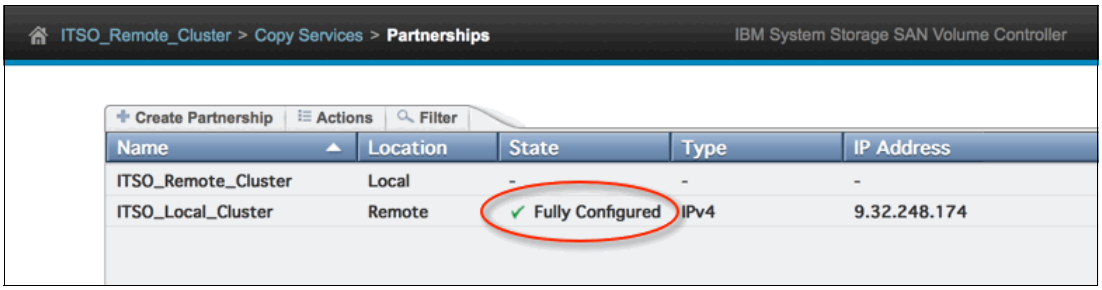

Figure 3-18 Fully configured IP partnership on ITSO_Remote_Cluster in the web-based GUI

The IP partnership is now fully configured, and you can proceed to configure the remote-copy relationships.

3.1.5 Creating IP remote-copy consistency groups and relationships

This section describes how to configure and activate remote-copy relationships that will use the IP replication partnership created in 3.1.4, “Create and configure the IP partnership on each of the partner systems” on page 28. Although IP replication supports all three types of remote-copy relationships (Metro Mirror (MM), Global Mirror (GM), and GM/CV), the advised remote-copy relationship type for use with IP replication is GM/CV. The following example shows relationships created using GM/CV.

Prerequisites and example environment details

To successfully create a remote-copy relationship, you must ensure that, for each volume that you want to replication, the remote partner system that you want to replicate to contains an available auxiliary volume that is the same size as the master volume. In this example, ITSO_Local_Cluster contains the following four 17 GB master volumes:

- repl_master_1
- repl_master_2
- repl_master_3
- repl_master_4

ITSO_Local_Cluster also contains the following four change volumes:

- repl_master_cv_1
- repl_master_cv_2
- repl_master_cv_3
- repl_master_cv_4

Additionally, ITSO_Remote_Cluster contains the following four 17 GB auxiliary volumes:

- repl_auxiliary_1
- repl_auxiliary_2
- repl_auxiliary_3
- repl_auxiliary_4

It also contains the following four change volumes:

- repl_auxiliary_cv_1
- repl_auxiliary_cv_2
- repl_auxiliary_cv_3
- repl_auxiliary_cv_4

The example uses a consistency group called *ITSO_Repl_CG* to contain the 4 volumes that we will replicate from ITSO_Local_Cluster to ITSO_Remote_Cluster.

Creating IP remote-copy consistency group and relationships on ITSO_Local_Cluster

To create a remote-copy relationship, perform the following steps:

1. Access the Remote Copy pane in the web-based GUI of ITSO_Local_Cluster using the **Copy Services** icon in the left navigation menu, as shown in Figure 3-19.

Figure 3-19 Accessing the Remote Copy pane in the web-based GUI

2. On the Remote Copy pane, click **New Consistency Group** at the top of the pane to create a consistency group containing multiple remote-copy relationships, as shown in Figure 3-20. Alternatively, if you choose not to use a consistency group, click **New Relationship** at the top of the pane to create a remote-copy relationship for a single volume.

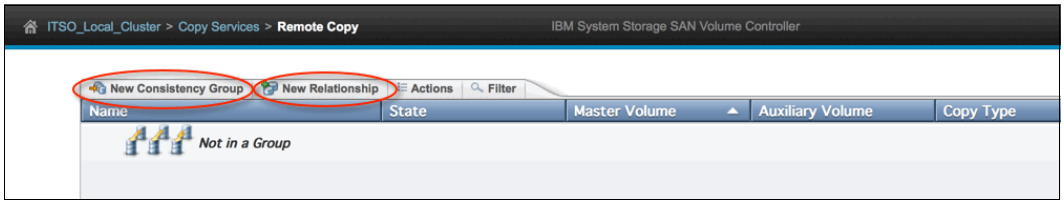

Figure 3-20 Creating a remote-copy relationship in the web-based GUI

3. In the New Consistency Group dialog window, enter the Consistency Group Name and click **Next**, as shown in Figure 3-21.

Figure 3-21 Creating a New Consistency Group for IP Replication in the web-based GUI

4. In the New Consistency Group dialog window, indicate the location of the auxiliary volumes and click **Next**, as shown in Figure 3-22. In this example, the auxiliary volumes are located on ITSO_Remote_Cluster, called *ITSO_Remote_Cluster*.

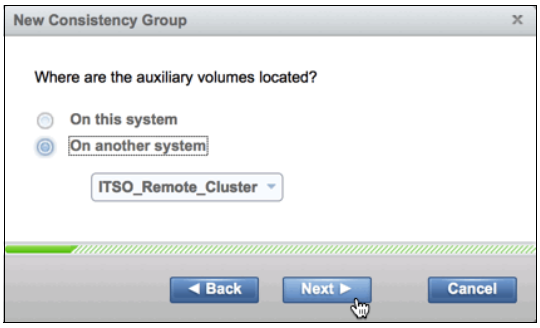

Figure 3-22 Indicating the location of the auxiliary volumes in the web-based GUI

5. In the New Consistency Group dialog window, indicate whether you want to add remote-copy relationships to the consistency group now or later and click **Next**, as shown in Figure 3-23. In this example, choose to add the relationships now.

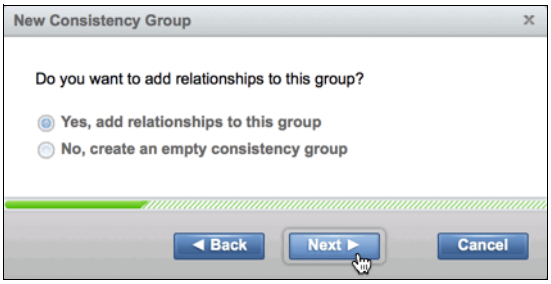

Figure 3-23 Indicating when to add remote-copy relationships to the consistency group in the web-based GUI

6. In the New Consistency Group dialog window, indicate the type of remote-copy relationship you want to add and click **Next**, as shown in Figure 3-24. In this example, choose **Global Mirror with Change Volumes**.

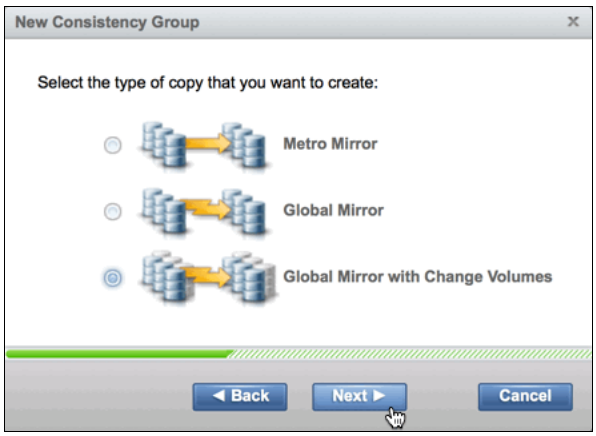

Figure 3-24 Choosing the type of remote-copy relationship in the web-based GUI

7. In the New Consistency Group dialog window, select any existing remote-copy relationships you want to add to this consistency group, as shown in Figure 3-25. In this example, choose to create new relationships, so click **Next**.

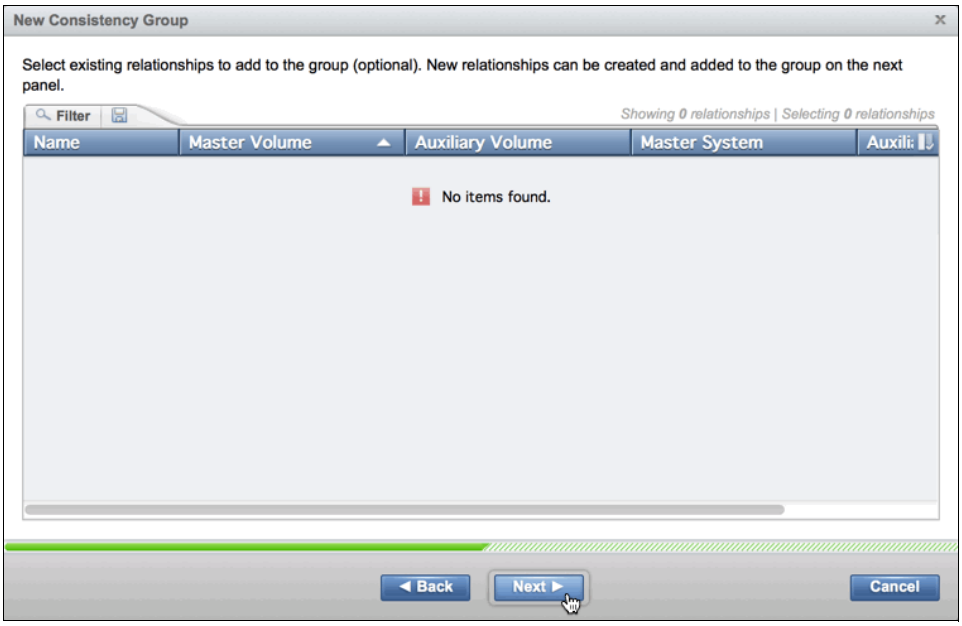

Figure 3-25 Selecting existing remote-copy relationships, if any, in the web-based GUI

8. In the New Consistency Group dialog window, use the drop-down menus to select the master and auxiliary volume relationship pairs, as shown in Figure 3-26.

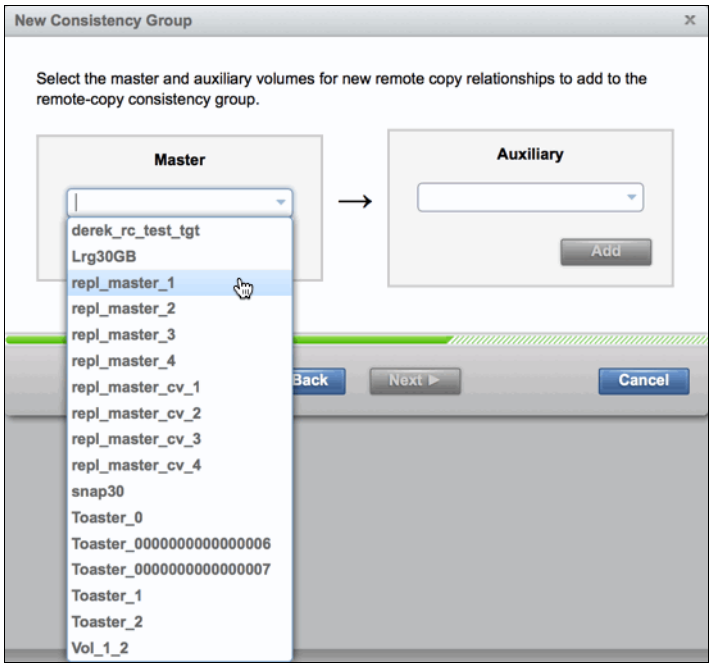

Figure 3-26 Selecting primary and auxiliary relationship volumes in the web-based GUI

9. Click **Add** to create a relationship between each selected pair. For each pair that you add, the Add New Global Mirror Change Volume dialog window displays, as shown in Figure 3-27.

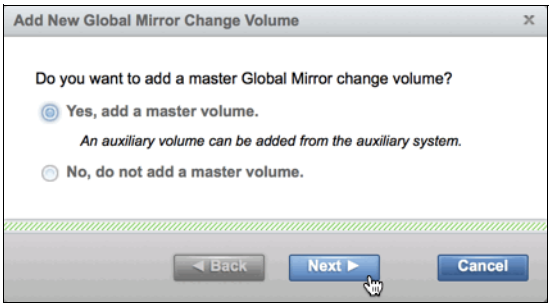

Figure 3-27 Indicating when to add master change volumes in the web-based GUI

10. You can choose to create or add the master Global Mirror change volume at this point, or after the consistency group has been fully created. In this example, choose to add the master change volumes at this step and click **Next**.

Important: Master change volumes can be created or added at this step, or at a later point. You *must* create or add auxiliary change volumes from the auxiliary system GUI or CLI after the remote-copy relationships have been added, or after the consistency group configuration is complete. The remote-copy relationship *cannot* be started until the change volumes have been created or added to the relationship for the master and auxiliary volumes.

11. In the Add New Global Mirror Change Volume dialog window, use the radio buttons to indicate whether you want to create a new master change volume, or use an existing master change volume, as shown in Figure 3-28. In this example, you indicate that you will use an existing master change volume, then use the drop-down menu to select it. Click **Finish** to complete adding this relationship to the consistency group.

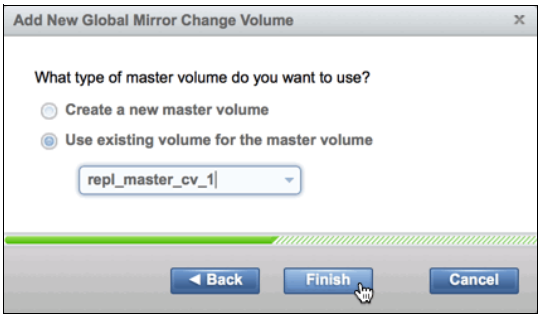

Figure 3-28 Selecting the master Global Mirror change volume in the web-based GUI

Figure 3-29 shows that the New Consistency Group dialog window updates to reflect the remote-copy relationship that you just added to the consistency group.

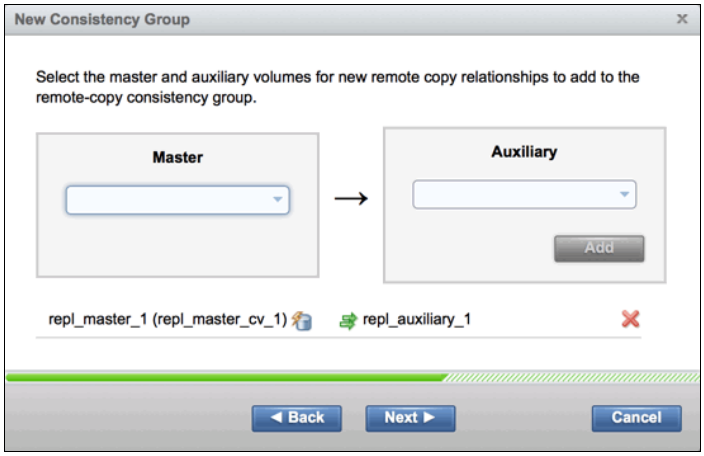

Figure 3-29 A new remote-copy relationship added to a new Consistency Group in the web-based GUI

12. Repeat steps 8 on page 36 through 11 on page 37 for any additional remote-copy master and auxiliary volume relationships that you want to add to the consistency group. In this example, add the remaining three master/auxiliary volume relationships, as shown in Figure 3-30, and click **Next**.

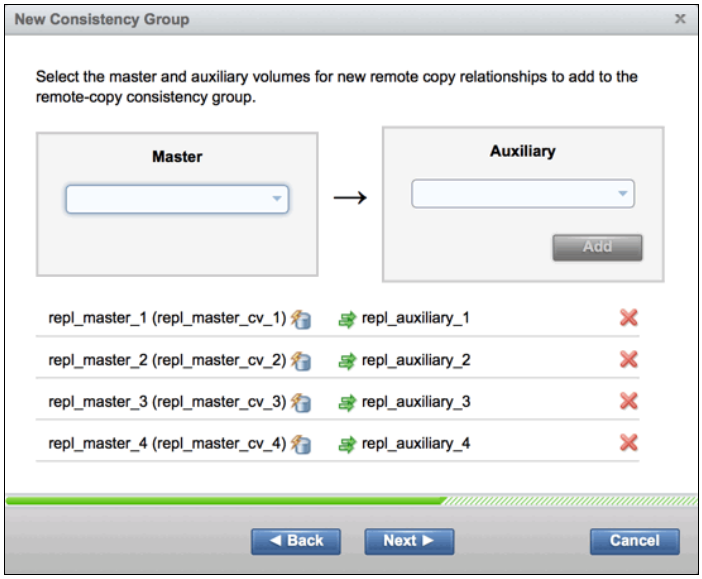

Figure 3-30 Four remote-copy relationships added to the consistency group in the web-based GUI

13. In the New Consistency Group dialog window, use the radio buttons to indicate whether the master and auxiliary volumes are already synchronized, as shown in Figure 3-31. In this example, indicate that the volumes are *not* synchronized, and then click **Next**.

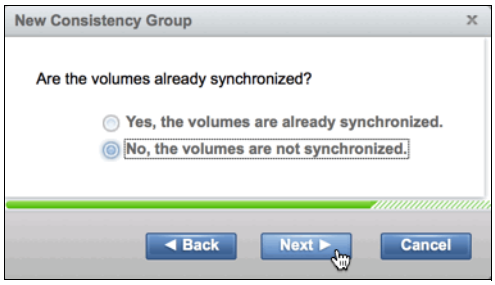

Figure 3-31 Indicating remote-copy volume synchronization state in the web-based GUI

14. In the New Consistency Group dialog window, use the radio buttons to indicate whether you want to start copying the volumes now or later, as shown in Figure 3-32. In this example, choose to use Global Mirror with Change Volumes. In this scenario, you must add the change volumes on the auxiliary system before you can start copying, so choose not to start copying and click **Finish**.

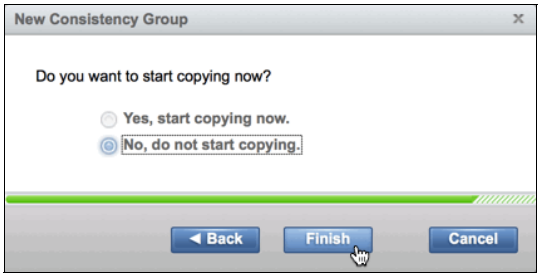

Figure 3-32 Selecting when to start copying the consistency group in the web-based GUI

15. The Create Remote-Copy Consistency Group dialog window displays the task status for creating the remote-copy consistency group, as shown in Figure 3-33. When the task completes, click **Close** to close the task status window.

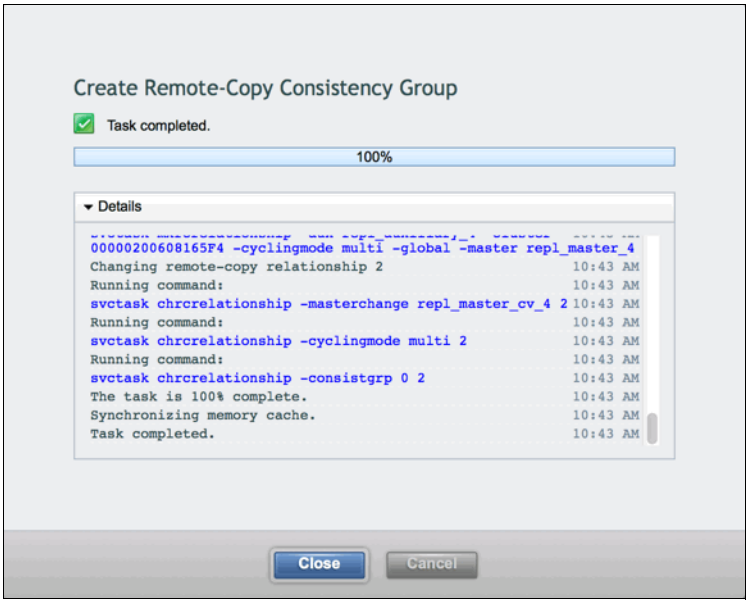

Figure 3-33 Task status for remote-copy consistency group creation in the web-based GUI

16. Confirm that the Remote Copy pane displays the consistency group and related remote-copy relationships that you just created, as shown in Figure 3-34.

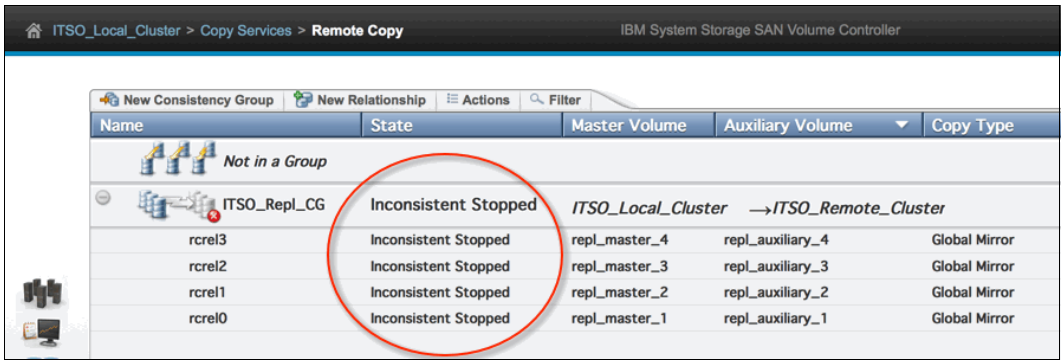

Figure 3-34 Verifying remote-copy consistency group creation in the web-based GUI

Note: Make note of the **State** of the consistency group and remote-copy relationships. In the example environment, we chose to use GM/CV. In this scenario, the consistency group *cannot* be started until you configure the change volumes on the auxiliary system. When creating consistency groups without change volumes, you might choose to start the copying as soon as the consistency group is created. In those cases, the value displayed for the **State** of the consistency group might differ.

For further information about consistency groups and remote copy status, see Chapter 2, “IP replication configuration overview and considerations” on page 11.

Completing IP remote-copy consistency group and relationship configuration on ITSO_Remote_Cluster

To complete the configuration, perform the following steps:

1. Access the Remote Copy pane in the web-based GUI of ITSO_Remote_Cluster using the **Copy Services** icon in the left navigation menu, as shown in Figure 3-35.

Figure 3-35 Accessing the Remote Copy pane in the web-based GUI

2. On the Remote Copy pane, you will see the same consistency group that was created on ITSO_Local_Cluster, as shown in Figure 3-36. If necessary, expand the consistency group to display the associated remote-copy relationships.

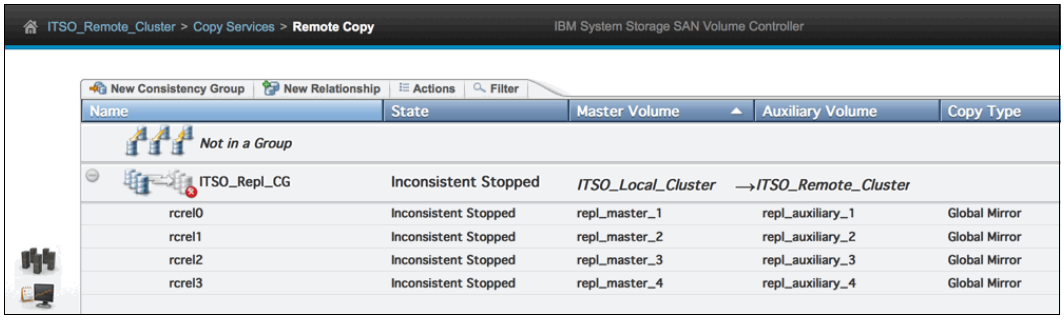

Figure 3-36 Verifying the Partner System consistency group in the web-based GUI

3. Right-click the row for the first remote-copy relationship and use the **Global Mirror Change Volumes** menu to add a change volume to this auxiliary volume. You can choose to create a new change volume or add an existing volume. In this example, choose to use an existing volume, so select **Global Mirror Change Volumes** → **Add Existing**, as shown in Figure 3-37.

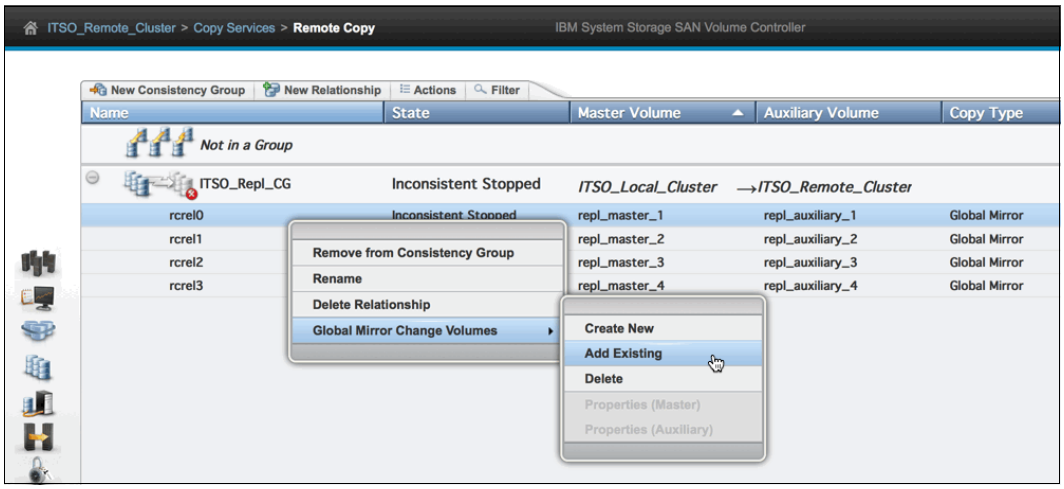

Figure 3-37 Adding auxiliary change volumes in the web-based GUI

4. In the Add Existing Global Mirror Change Volume dialog window, use the drop-down menu to select the appropriate auxiliary change volume, and click **Add**, as shown in Figure 3-38.

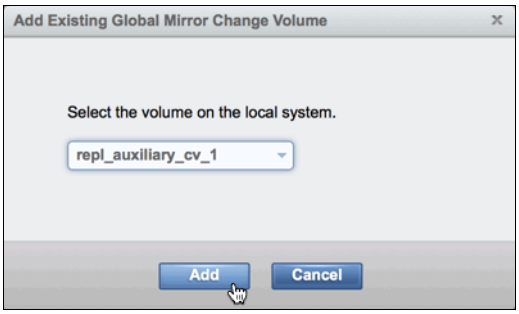

Figure 3-38 Selecting the auxiliary change volume in the web-based GUI

5. The Add Global Mirror Change Volume dialog window displays the task status for adding the change volume, as shown in Figure 3-39. When the task completes, click **Close** to close the task status window.

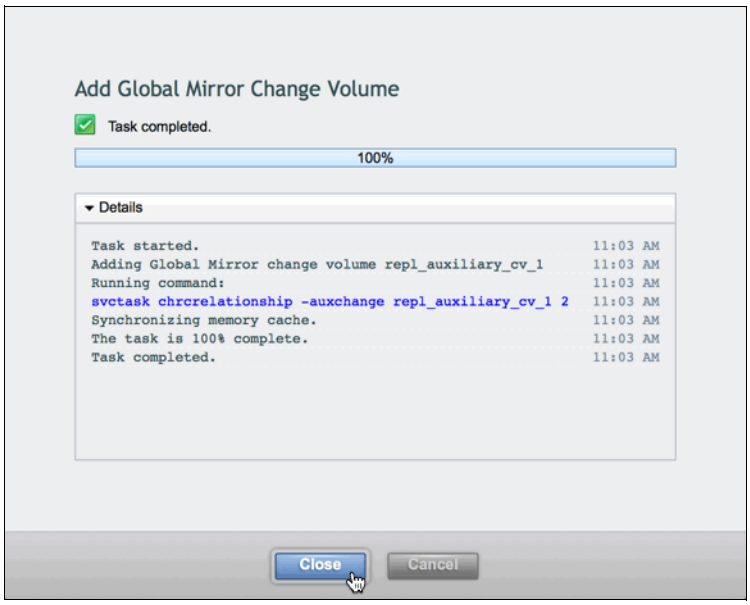

Figure 3-39 Task status for adding a change volume in the web-based GUI

6. Repeat steps 3 on page 41 and 4 on page 41 for each of the remaining remote-copy relationships in the consistency group.

Starting the IP remote-copy consistency group

To start the IP remote-copy consistency group, perform the following steps:

1. After all of the auxiliary change volumes have been added to the consistency group, right-click the consistency group in the Remote Copy pane and select **Start** to start the consistency group, as shown in Figure 3-40.

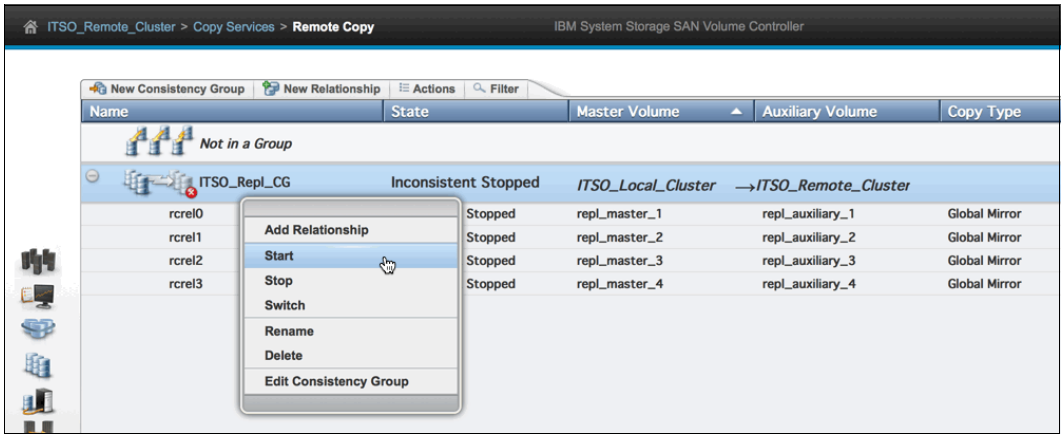

Figure 3-40 Starting the IP replication remote-copy consistency group in the web-based GUI

2. The Start Remote-Copy Consistency Group dialog window displays the task status for starting the consistency group, as shown in Figure 3-41 on page 43. When the task completes, click **Close** to close the task status window.

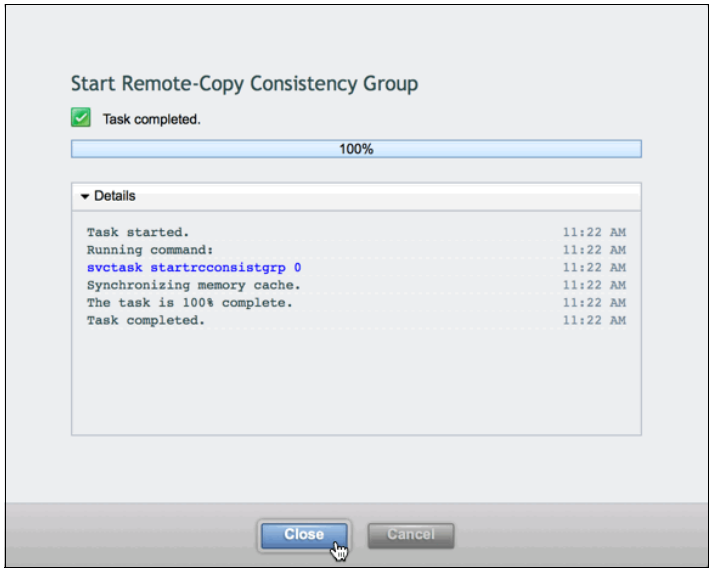

Figure 3-41 Task status for starting a remote-copy consistency group in the web-based GUI

Verifying and monitoring IP remote-copy status

To verify and monitor the IP remote-copy status, perform the following steps:

1. The Remote Copy pane should indicate that the consistency group state is now *Inconsistent Copying*, as shown in Figure 3-42.

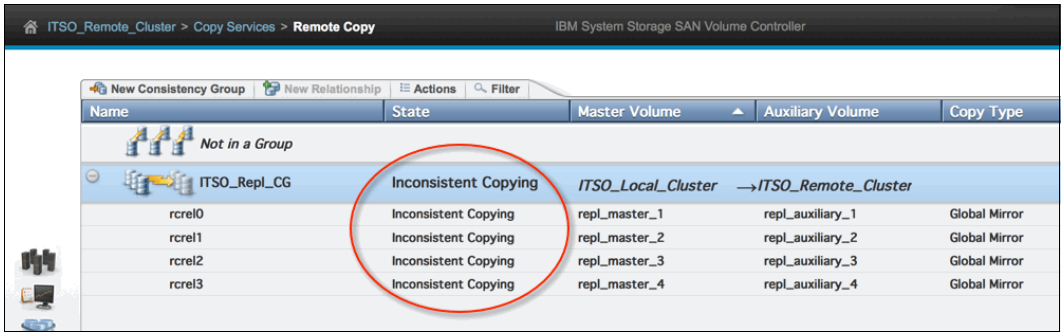

Figure 3-42 Verifying remote-copy consistency group state in the web-based GUI

2. Additionally, you can verify that the expected Ethernet ports are in use for IP replication by clicking **Settings** → **Network** from the left navigation menu, then accessing the **Network** → **Ethernet Ports** pane, as shown in Figure 3-43. Note the value in the **Remote Copy Active** column.

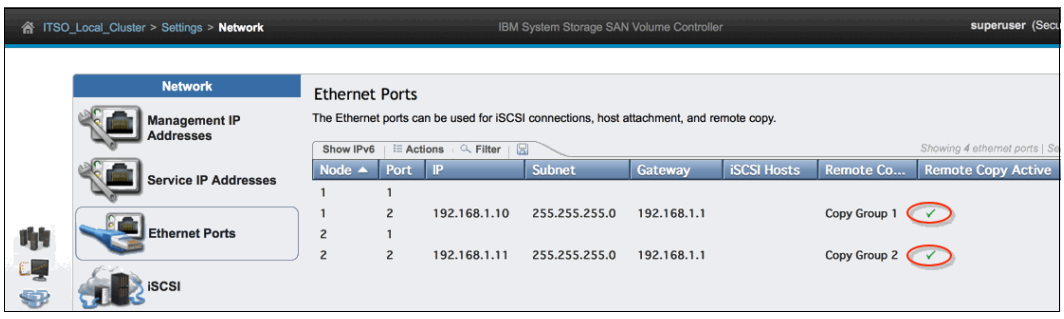

Figure 3-43 The Ethernet Ports pane

3. Further, you can monitor the IP remote-copy bandwidth usage of the system from the Performance page. From the left navigation menu, choose **Monitoring** → **Performance** menu, as shown in Figure 3-44.

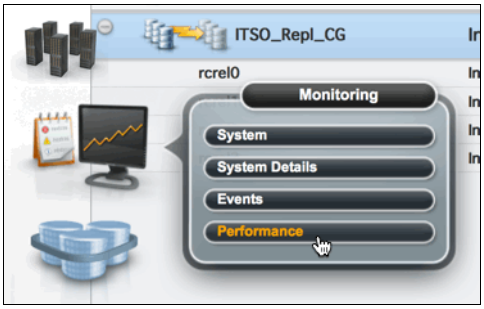

Figure 3-44 Accessing the Performance page in the web-based GUI

4. In the **Interfaces** graph in the lower left region of the **Performance** pane, ensure that the **IP Remote Copy** box is selected, as shown on Figure 3-45. When selected, the graph will display IP Remote Copy statistics.

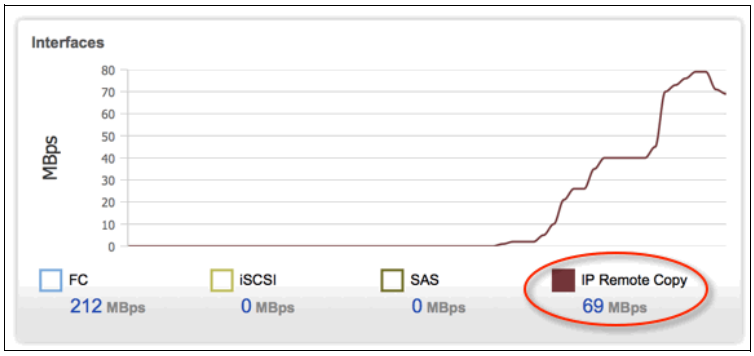

Figure 3-45 Monitoring IP Remote Copy performance in the web-based GUI

4

Best practices and troubleshooting

This chapter captures some best practices for Internet Protocol (IP) replication based on field experience. It also provides an overview of common problems that can occur during the implementation and usage of IP replication.

This chapter includes the following sections:

- Best practices
- Implementation and usage troubleshooting

4.1 Best practices

This section describes best practices based on field experience of IP replication.

4.1.1 System configuration

IP replication provides for the use of single or dual input/output (I/O) groups, in addition to single or dual remote-copy port groups. By using one or the other, speed *or* redundancy can be achieved for replication. By using both, speed *and* redundancy can be achieved.

Figure 4-1 shows an optimal configuration using both dual I/O Groups and dual remote copy port groups.

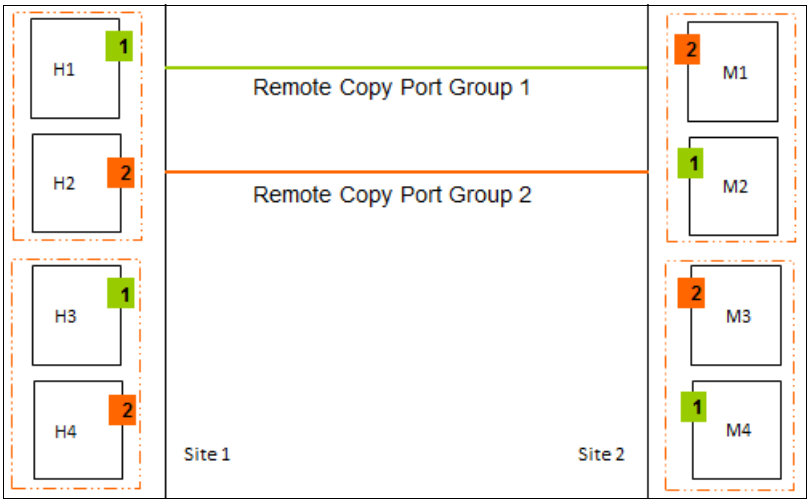

Figure 4-1 Dual I/O group and dual remote copy port group example

Restriction: IBM Storwize V7000 and IBM System Storage SAN Volume Controller (SVC) nodes with more than two I/O groups are supported. However, there is a limitation that only four IP addresses from two I/O groups can be included in a remote copy port group.

4.1.2 Performance considerations

The following advice has come out of field experience:

- 2145-CG8 with request for price quotation (RPQ) 8S1296 is advised if IP replication and IBM Real-time Compression are being used. This is advised so that processor-intensive services have adequate resources for optimal usability.
- If the nodes performing IP replication are 2145-CF8 or earlier, *and* Real-time Compression is being used, configure ports for the IP replication traffic on I/O groups not containing compressed volumes. This is advised so that processor-intensive services have adequate resources for optimal usability.
- If Internet Small Computer System Interface (iSCSI) is used for host connectivity, do not use the same port for host I/O and IP replication traffic. Optimal performance can be achieved by isolating the different traffic onto different ports.

Important: Decreased throughput on the source back-end controller can adversely affect the consistency of the throughput of IP replication traffic.

4.2 Implementation and usage troubleshooting

This section describes some of the issues that might arise during the initial implementation of IP replication.

4.2.1 Error code CMMVC8359E

During the implementation of IP replication, error code CMMVC8359E might be encountered.

Figure 4-2 shows the specific error as it is encountered in the graphical user interface (GUI).

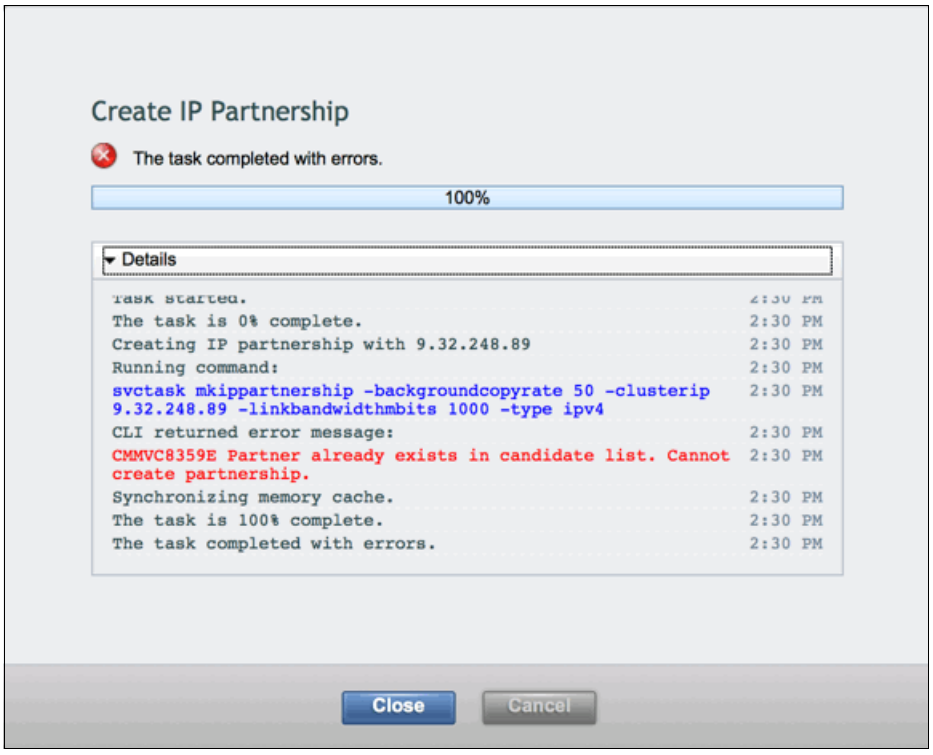

Figure 4-2 CMMVC8359E error pane

This error is caused by the destination system being in a Fibre Channel (FC) zone with the source system. Ensure that the two systems are not zoned together, and attempt to create the partnership again.

4.2.2 Error code CMMVC8354E

During the implementation of IP replication, error code CMMVC8354E might be encountered.

Figure 4-3 shows the specific error as it is encountered in the GUI.

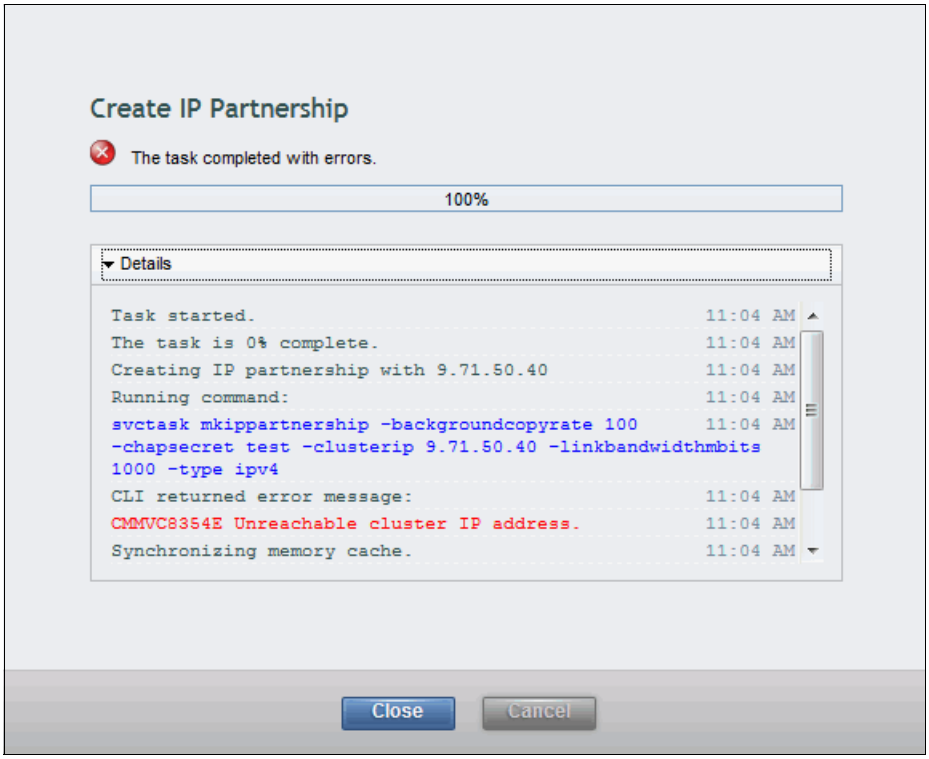

Figure 4-3 CMMVC8354E error pane

There are three common reasons why this error can occur:

- A lack of communication between the source system and the destination system over port 3260. Ensure that Transmission Control Protocol (TCP) port 3260 is open bidirectionally between the source and destination systems.
- The source and destination system are not in the same layer.
- The IP address input into the partnership is incorrect.

4.2.3 Error Code CMMVC8353E

During implementation and usage of IP replication, error code CMMVC8353E might be encountered.

Figure 4-4 shows the specific error as it is encountered in the GUI.

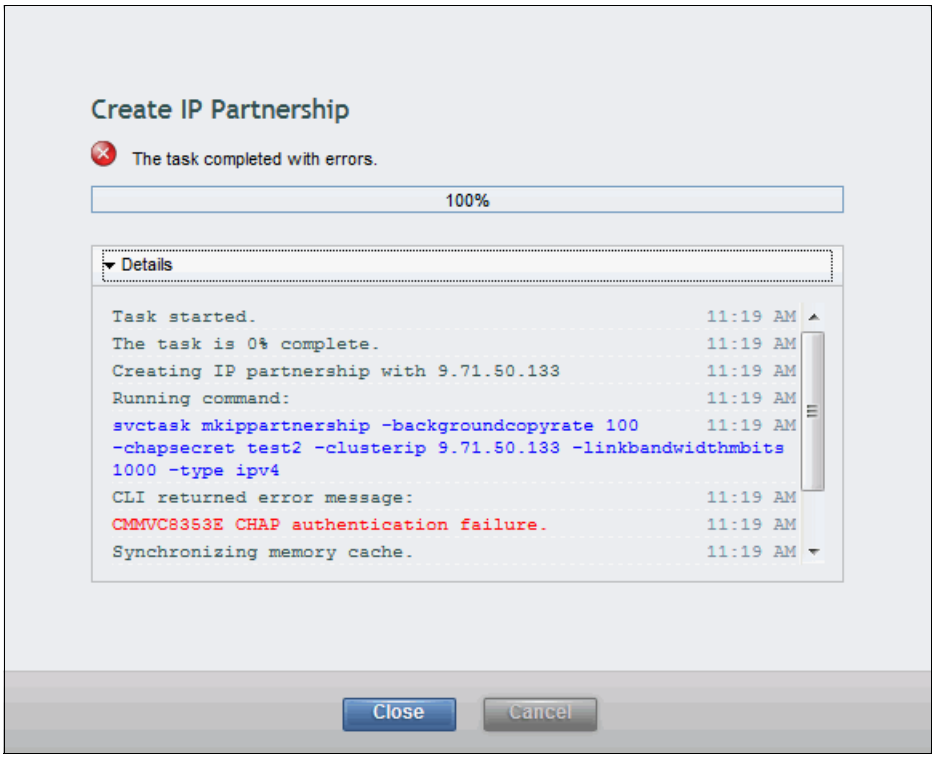

Figure 4-4 CMMVC8353E error pane

This error is caused by the Challenge Handshake Authentication Protocol (CHAP) secret setup in the partnership not matching the CHAP secret defined in the destination system. Ensure that the CHAP secret of the partnership matches the CHAP secret of the destination system. For more information about CHAP secret, see Appendix B, “CHAP Information and Implementation” on page 57.

4.2.4 Error Code 2020

This error can occur both during implementation and during normal usage of IP Replication. This error occurs when the Remote Copy link is unavailable.

Figure 4-5 shows the Event Log when an error code 2020 appears.

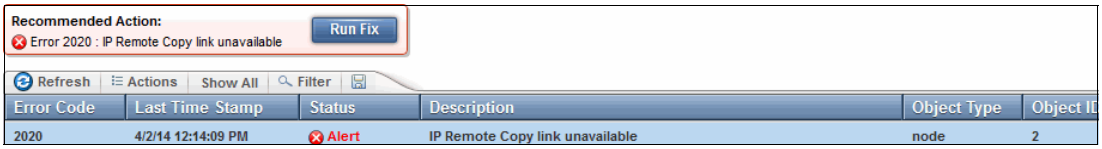

Figure 4-5 Event Log showing error 2020

Figure 4-6 shows the Directed Maintenance Procedure (DMP) pane for the error. This pane shows general information about what to check in the LAN environment to assist with diagnosing this error.

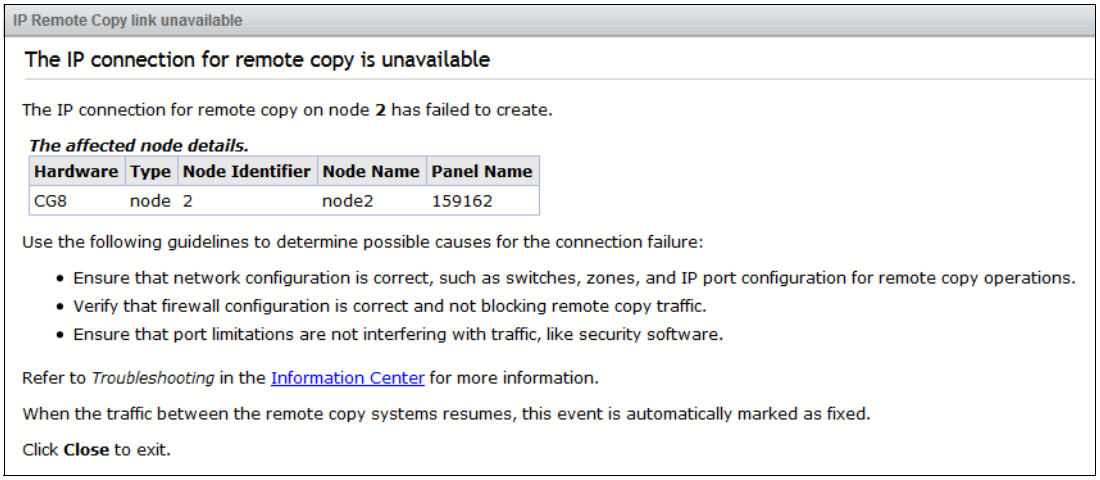

Figure 4-6 DMP information pane for the 2020 error code

There are multiple reasons why this error can occur:

- A firewall or switch has not been properly configured to enable communication over port 3265.
- An external system is blocking access to port 3265.
- Other LAN issues exist (for example, cabling, switches, and so on).

4.2.5 Error Code 2021

This error can occur during normal usage of IP replication. This error occurs when the destination system IP is not reachable.

Figure 4-7 shows the Event Log when an error code 2021 appears.

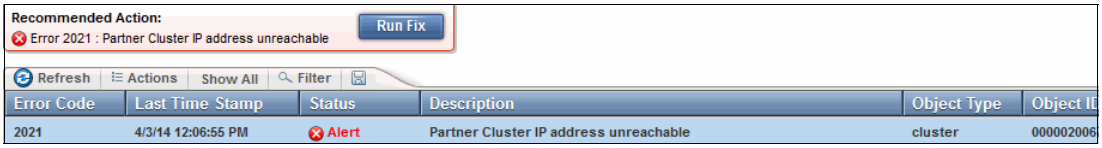

Figure 4-7 Event Log showing error 2021

Figure 4-8 shows the DMP pane for the error. It has information regarding the partner system that is affected. From here, you can check the connectivity of the destination system, or modify the address of the destination system in the partnership, if required.

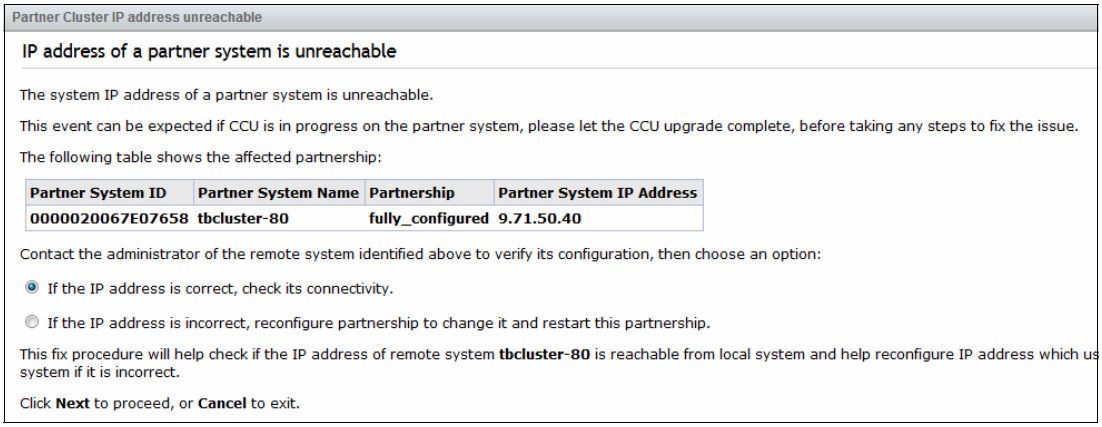

Figure 4-8 DMP information pane for the 2021 error code

There are multiple reasons why this error can occur:

- The IP of the destination system has been modified, and the partnership information has not been updated.
- A firewall or switch has not been properly configured to enable communication over port 3260.
- An external system is blocking access to port 3260.
- Other LAN issues exist (for example, cabling, switches, and so on).

4.2.6 Error Code 2022

This error can occur during normal usage of IP replication. This error occurs when the CHAP secret configured in the IP partnership does not match the system CHAP secret of the remote system.

Figure 4-9 shows the Event Log when an error code 2022 appears.

Figure 4-9 Event Log showing error 2022

This error can occur when the system CHAP secret is modified on one partner system, and the IP partnership CHAP secret is not modified appropriately on the partner system.

Figure 4-10 shows the DMP pane for the error. It has information regarding the partner system that is disallowing CHAP authentication.

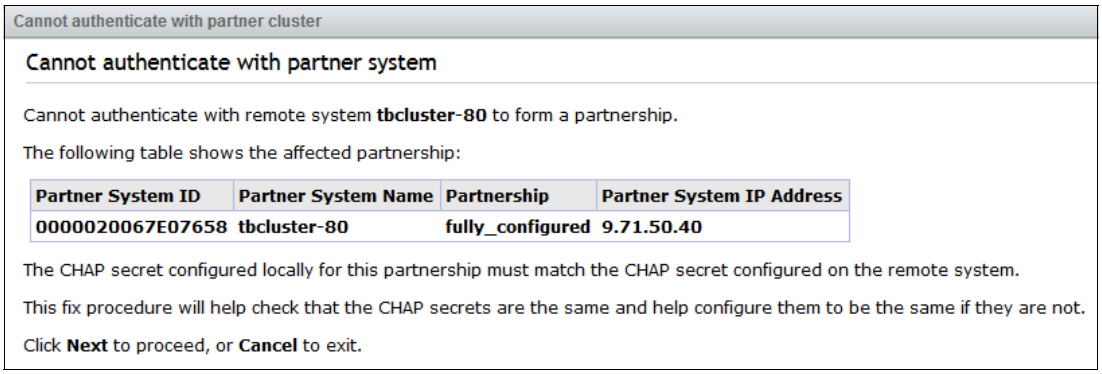

Figure 4-10 DMP Information Pane for 2022 error code

Figure 4-11 shows the DMP pane that enables you to modify the partnership CHAP secret to match the system CHAP secret of the destination system.

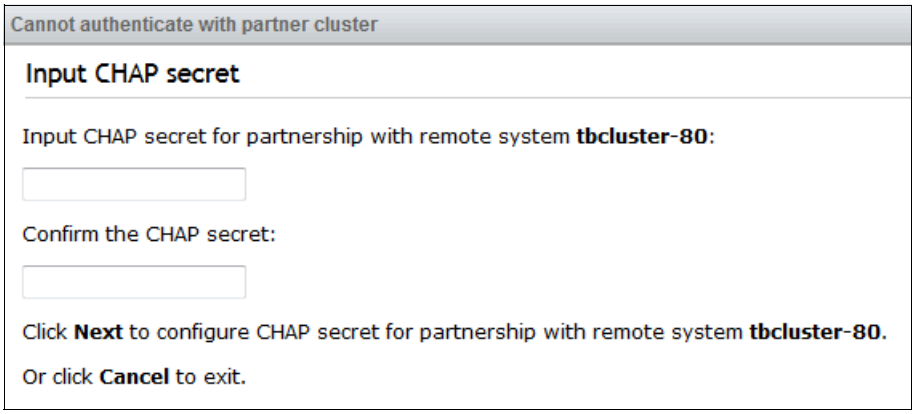

Figure 4-11 Additional DMP information pane for 2022 error code

This error should be fixed so that if the partnership is stopped for whatever reason, it will be able to authenticate with the destination system after the partnership starts again.

4.2.7 Error Code 2023

This error can occur during normal usage of IP replication. This error occurs when the system ID of the partner system has been changed, or is not expected.

These are the main reasons why this error can occur:

- The IP address of the partner system has changed, and the IP address configured in the partnership is pointing to a different system.
- The partner system has experienced a system recovery, and the unique system ID has been changed.

To resolve this error, the partnership needs to be removed and re-created.

A

Command-line interface

The following outlines the command-line interface (CLI) commands required to complete the operations carried out in Chapter 2, “IP replication configuration overview and considerations” on page 11, using the GUI, to configure IP replication between two IBM System Storage SAN Volume Controller (SVC) systems. Each SVC system consists of a single input/output (I/O) group using port 2 on each node. Two intersite links are used.

For CLI commands used in configuring Challenge Handshake Authentication Protocol (CHAP) see Appendix B, “CHAP Information and Implementation” on page 57.

Configuring Ethernet ports on the partner systems

To configure the Ethernet port of each node to be used for IP replication, perform the following steps:

1. Run the **cfgportip** command shown in Example A-1.

Example A-1 Configure network ports

```
IBM_2145:ITSO_Local_Cluster:superuser>cfgportip -gw 192.168.1.1 -host no -ip
192.168.1.10 -mask 255.255.255.0 -node 1 -remotecopy 1 2
IBM_2145:ITSO_Local_Cluster:superuser>
```

2. Confirm changes using the **lsportip** command shown in Example A-2.

Example A-2 The lsportip command output

```
IBM_2145:ITSO_Local_Cluster:superuser>lsportip
id node_id node_name IP_address   mask          gateway     IP_address_6
prefix_6 gateway_6 MAC               duplex state        speed failover
link_state host remote_copy host_6 remote_copy_6 remote_copy_status
remote_copy_status_6
1  1      node1
e4:1f:13:30:19:e4 Full   unconfigured 1Gb/s no      active        0
0
1  1      node1
e4:1f:13:30:19:e4 Full   unconfigured 1Gb/s yes     active        0
0
2  1     node1    192.168.1.10 255.255.255.0 192.168.1.1
e4:1f:13:30:19:e6 Full   configured   1Gb/s no      active    no  1
0             unused
2  1      node1
e4:1f:13:30:19:e6 Full   configured   1Gb/s yes     active        0
0
1  2      node2
e4:1f:13:b7:7f:78 Full   unconfigured 1Gb/s no      active        0
0
1  2      node2
e4:1f:13:b7:7f:78 Full   unconfigured 1Gb/s yes     active        0
0
2  2     node2    192.168.1.11 255.255.255.0 192.168.1.1
e4:1f:13:b7:7f:7a Full   configured   1Gb/s no      active    no  2
0             unused
2  2      node2
e4:1f:13:b7:7f:7a Full   configured   1Gb/s yes     active        0
0
IBM_2145:ITSO_Local_Cluster:superuser>
```

3. This example shows only one port being configured. Run the command on *all* ports to be configured on *each* partner system.

   In this example, we chose not to use the port as an Internet Small Computer System Interface (iSCSI) host, and as a member of a port group for replication. Although this is supported, it is not recommended. See 2.1.2, “IP partnership requirements” on page 13 for more details.

Create and configure the IP partnership on each of the partner systems

This section describes how to create and configure the IP partnership:

1. Run the **mkippartnership** command to configure the creation of the IP partnership on each system, as shown in Example A-3.

Example A-3 Configuring IP partnerships

```
IBM_2145:ITSO_Local_Cluster:superuser>mkippartnership -backgroundcopyrate 50
-clusterip 9.32.248.89 -linkbandwidthmbits 2000 -type ipv4
IBM_2145:ITSO_Local_Cluster:superuser>
```

2. In this example, we set the link bandwidth to 2000 (2 x 1000 megabits per second (Mbps) links) with a background copy rate of 50%. Do this on both systems in the partnership. Initially, the partnership will show the `partially_configured_local` status output, as shown in Example A-4.

Example A-4 The lspartnership command output for the first system defined

```
IBM_2145:ITSO_Local_Cluster:superuser>lspartnership
id               name                location partnership                type
cluster_ip  event_log_sequence
0000020060C14FE2 ITSO_Local_Cluster  local
00000200608165F4 ITSO_Remote_Cluster remote   partially_configured_local ipv4
9.32.248.89
IBM_2145:ITSO_Local_Cluster:superuser>
```

3. Run the same commands on the remote partner system. After the partnership has been created on the second system, the partnership will change to `fully configured`.

Remember: Both candidates for the partnership need to be in the same replication layer. When creating the partnership, use the management IP address of the remote system as the clusterip.

Creating IP remote-copy consistency groups and relationships

Because the advised usage of IP replication is with GM/CV, this is what we use in this example. However, Metro Mirror (MM) and Global Mirror (GM) are supported. For any mirror relationship to be created, a volume of the same size needs to exist on both systems in the partnership. This example used four volumes in a consistency group, but for clarity only the creation of the first volume's remote copy relationship is shown in Example A-5.

Example A-5 Creating consistency group and mirror relationships

```
IBM_2145:ITSO_Local_Cluster:superuser>mkrcconsistgrp -cluster 00000200608165F4
-name ITSO_repl_CG
IBM_2145:ITSO_Local_Cluster:superuser>
IBM_2145:ITSO_Local_Cluster:superuser>mkrcrelationship -aux repl_auxiliary_1
-cluster 00000200608165F4 -global -master repl_master_1
IBM_2145:ITSO_Local_Cluster:superuser>
IBM_2145:ITSO_Local_Cluster:superuser>svctask chrcrelationship -masterchange
repl_master_cv_1 3
IBM_2145:ITSO_Local_Cluster:superuser>
IBM_2145:ITSO_Local_Cluster:superuser>svctask chrcrelationship -cyclingmode multi
3
```

```
IBM_2145:ITSO_Local_Cluster:superuser>
IBM_2145:ITSO_Local_Cluster:superuser>svctask chrcrelationship -consistgrp 0 3
IBM_2145:ITSO_Local_Cluster:superuser>
```

When all of the volume relationships are created, they will show as **in**consistent stopped, as shown in Example A-6.

Example A-6 The lsrcrelationship command output

```
IBM_2145:ITSO_Local_Cluster:superuser>lsrcrelationship
id name   master_cluster_id master_cluster_name master_vdisk_id master_vdisk_name
aux_cluster_id   aux_cluster_name    aux_vdisk_id aux_vdisk_name   primary
consistency_group_id consistency_group_name state                bg_copy_priority
progress copy_type cycling_mode freeze_time
2  rcrel3 0000020060C14FE2  ITSO_Local_Cluster  2               repl_master_4
00000200608165F4 ITSO_Remote_Cluster 6       repl_auxiliary_4 master  0
ITSO_repl_CG           inconsistent_stopped 50               0        global
multi
3  rcrel0 0000020060C14FE2  ITSO_Local_Cluster  3               repl_master_1
00000200608165F4 ITSO_Remote_Cluster 2       repl_auxiliary_1 master  0
ITSO_repl_CG           inconsistent_stopped 50               0        global
multi
6  rcrel1 0000020060C14FE2  ITSO_Local_Cluster  6               repl_master_2
00000200608165F4 ITSO_Remote_Cluster 3       repl_auxiliary_2 master  0
ITSO_repl_CG           inconsistent_stopped 50               0        global
multi
7  rcrel2 0000020060C14FE2  ITSO_Local_Cluster  7               repl_master_3
00000200608165F4 ITSO_Remote_Cluster 5       repl_auxiliary_3 master  0
ITSO_repl_CG           inconsistent_stopped 50               0        global
multi
IBM_2145:ITSO_Local_Cluster:superuser>
```

When the volume relationships are all created and the volumes have their associated change volume assigned on the local system, you will need to add the change volumes to the mirrored volumes on the partner/remote system. In this example, it is ITSO_Remote_Cluster, as shown in Example A-7. For clarity, this is shown on the first volume only.

Example A-7 Adding change volumes to remote partner system

```
IBM_2145:ITSO_Remote_Cluster:superuser>chrcrelationship -auxchange
repl_auxiliary_cv_1 2
IBM_2145:ITSO_Remote_Cluster:superuser>
```

Starting the IP remote-copy consistency group

When all of the auxiliary change volumes have been added, the final step is to start the remote-copy consistency group, as shown in Example A-8.

Example A-8 Starting the remote copy

```
IBM_2145:ITSO_Local_Cluster:superuser>
IBM_2145:ITSO_Local_Cluster:superuser>startrcconsistgrp 0
IBM_2145:ITSO_Local_Cluster:superuser>
```

IP Replication is now configured.

B

CHAP Information and Implementation

This appendix provides additional information about Challenge Handshake Authentication Protocol (CHAP) and how to properly implement it for use with Internet Protocol (IP) replication.

General information on CHAP

CHAP is used in many different situations and products, and for the IBM System Storage SAN Volume Controller (SVC)/IBM Storwize family, it is used for both IP replication and Internet Small Computer System Interface (iSCSI)-attached host authentication.

CHAP works by first creating a shared secret on each system. Then, during the link establishment phase, one side sends a challenge to the other. The second side then combines the challenge and the secret and hashes it. After hashing, the second side sends the information back to the first side, which compares it against its own calculated value. If the calculation matches the sent value, the connection is completed. Otherwise, the connection is rejected.

The next section will go into more detail regarding the specific setup of CHAP with the SVC/Storwize family of products.

Configuring CHAP for IP replication

There are only a few steps that are needed to properly configure CHAP with IP replication:

1. Modify the system-wide CHAP secret for both the source and destination systems.
2. Configure the partnership from the source to the destination to include the CHAP secret of the destination system.
3. Configure the partnership from the destination to the source to include the CHAP secret of the source system.

Tip: The CHAP secret used on each system does not have to be the same. The CHAP secret field in the partnership must match the CHAP secret of the system that it is connecting to.

Configure system-wide CHAP secret

Using the graphical user interface (GUI), there are many different paths to get to the Modify CHAP Configuration pane, which is needed to set up a system-wide CHAP secret for use with IP Replication. You can select any one of the following choices:

- **Copy Services** → **Partnerships** → **Actions** → **Modify CHAP Configuration**
- **Monitoring** → **System Details** → **Actions** → **Modify CHAP Configuration**
- **Settings** → **Network** → **iSCSI** → **Modify CHAP Configuration**

After the Modify CHAP Configuration pane is opened, follow these steps:

1. Input the CHAP secret for the current system, select the box marked **Use for IP partnerships**, and press **Modify**.

Note 1: If you are modifying the CHAP secret on a system that is running an IP replication partnership, the partnership must be in a `stopped` state before modifying.

Note 2: If you are using the GUI to remove the CHAP secret on a system, clear the **Use for IP partnerships** box first, then remove the password. Click **Modify** to proceed.

2. Figure B-1 shows the pane that you should see when the system-wide CHAP is modified.

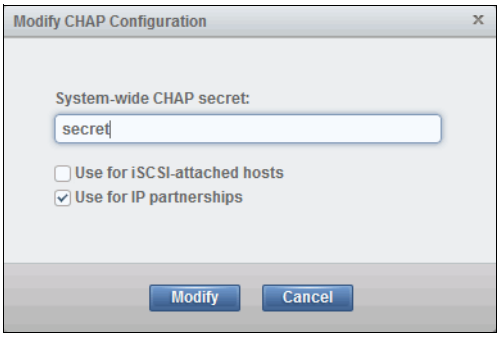

Figure B-1 Modify CHAP configuration

3. Repeat the previous procedure on the secondary system.

Information: Alternatively, you can use the following CLI command:

```
chsystem -chapsecret chap_secret
```

To remove a CHAP secret, you can use the following CLI command:

```
chsystem -nochapsecret
```

Configure partnership to use the CHAP secret

The CHAP secret can be added to a partnership either during initial creation, or during normal usage.

To do so, follow the steps in Chapter 3, “Implementing IP replication in the SVC/Storwize family” on page 21.

The CHAP secret of the destination system will be input into the **Partner system’s CHAP secret** field.

Alternatively, the partnership can be modified at any time by going to its properties, and adding or modifying the partner system’s CHAP secret.

Information: To make a partnership and include a CHAP secret using the CLI, use the following command:

```
mkippartnership -backgroundcopyrate 100 -chapsecret test -clusterip 9.71.50.133
-linkbandwidthmbits 1000 -type ipv4
```

To modify an existing partnership to include a CHAP secret, stop the partnership and use the following command:

```
chpartnership -chapsecret chap_secret clusterid
```

Related publications

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

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks”. Note that some of the documents referenced here might be available in softcopy only:

- *ABCs of z/OS System Programming Volume 1*, SG24-6981
- *ABCs of z/OS System Programming Volume 2*, SG24-6982
- *ABCs of z/OS System Programming Volume 3*, SG24-6983
- *ABCs of z/OS System Programming: Volume 4*, SG24-6984
- *ABCs of z/OS System Programming: Volume 5*, SG24-6985
- *ABCs of z/OS System Programming Volume 6*, SG24-6986
- *ABCs of z/OS System Programming Volume 7*, SG24-6987
- *ABCs of z/OS System Programming Volume 8*, SG24-6988
- *ABCs of z/OS System Programming: Volume 9*, SG24-6989
- *ABCs of z/OS System Programming Volume 10*, SG24-6990
- *ABCs of z/OS System Programming Volume 11*, SG24-6327
- *ABCs of z/OS System Programming Volume 12*, SG24-7621
- *ABCs of z/OS System Programming Volume 13*, SG24-7717
- *Implementing the IBM System Storage SAN Volume Controller V7.2*, SG24-7933
- *Implementing the IBM Storwize V7000 V7.2*, SG24-7938
- *Implementing the IBM Storwize V5000*, SG24-8162
- *Implementing the IBM Storwize V3700*, SG24-8107
- *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574

How to get IBM Redbooks

You can search for, view, or download IBM Redbooks, IBM Redpapers, Technotes, draft publications, and Additional materials, and order hardcopy Redbooks, at this website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads:

ibm.com/support

IBM Global Services:

ibm.com/services

IBM SAN Volume Controller and Storwize Family Native IP Replication

Implement SVC/Storwize Family native IP replication

Understand requirements for configuring native IP replication

Mitigate the effects of IP latency over distance

IBM has announced native Internet Protocol (IP) replication using Bridgeworks SANSlide technology with its IBM System Storage SAN Volume Controller (SVC), IBM Storwize V7000, IBM Storwize V5000 and Storwize V3700 virtualized storage systems.

This combination of SANSlide and the SVC/Storwize family provides a powerful solution for clients who require efficient, IP-based replication over long distances.

This certification gives SVC/Storwize clients a fully supported, transparent technology that includes unmatched levels of performance and reliability. With the SANSlide protocol acceleration technology, it is now possible to replicate data across continents in a cost-efficient way, with little or no loss in performance. At the same time, bandwidth usage can improve to over 95%, rather than the 1% - 5% normally achieved in long-distance IP networks.

This IBM Redpaper publication shows the steps required to implement this solution efficiently and speedily.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, clients, and IBM Business Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

REDP-5103-00