



Sangam Racherla
Alexander Watson

IBM System Storage DS3950 Introduction

Introduction/Overview

This IBM® Redpaper™ publication provides you with an overview of the new IBM System Storage® DS3950 offering and its functionality. Descriptions of the hardware components and the various cabling options are provided.

Because the DS3950 is identical in design and fit to the DS5020, it is useful to review the following related Redbooks® to obtain detailed information about this subsystem and its implementation procedures.

- ▶ *IBM Midrange System Storage Hardware Guide*, SG24-7676
- ▶ *IBM Midrange System Storage Implementation and Best Practices Guide*, SG24-6363
- ▶ *IBM Midrange System Storage Copy Services Guide*, SG24-7822

DS3950 Storage Subsystem

The new DS3950 storage subsystem is a version of the DS5020 storage subsystem which is offered in the European and some Asian Pacific geographies. This subsystem has most of the same capabilities as the DS5020. The major difference is the number of available host fiber channels attachments it can have.

The IBM System Storage DS3950 disk system, shown in Figure 1 on page 2, is designed to deliver high performance, advanced functionality, high availability, and modular and scalable storage capacity.




Figure 1 IBM System Storage DS3950 storage subsystem

The DS3950 is offered with the following configurations:

DS3950 disk system members

- ▶ DS3950 Disk Controller (1814-94H).
- ▶ DS3950 Disk Controller (1814-98H).
- ▶ EXP395 Storage Expansion (1814-92H).

Disk drives

- ▶ Up to 224 TB physical storage capacity.
- ▶ Accommodates up to 16 disk drives installed within the DS3950 enclosure.
- ▶ Attachment support for up to six EXP395 expansion enclosures.
- ▶ Attachment support for EXP810 with the Attach EXP810 to DS3950 Activation feature.
- ▶ Supports an intermix of Serial Advanced Technology Attachment (SATA) drives and FC drives within the DS3950 and EXP395 enclosures.

Host attachment

- ▶ DS3950 Disk Controller (1814-94H) - Four 8 GBps Fibre Channel (FC) ports for host connectivity.
- ▶ DS3950 Disk Controller (1814-98H) - Four 8 GBps FC and Four 1GBps iSCSI ports for host connectivity.

System cache

- ▶ DS3950 Disk Controller (1814-94H) - 2 GB cache memory (1 GB per RAID controller)
- ▶ DS3950 Disk Controller (1814-98H) - 4 GB cache memory (2 GB per RAID controller)

Drive options

- ▶ FC disks without encryption:
 - 300 GB/15K 4 Gbps FC E-DDM
 - 450 GB/15K 4 Gbps FC E-DDM
 - 600 GB/15K 4 Gbps FC E-DDM
- ▶ SATA disks:
 - 1000 GB/7.2K SATA E-DDM
 - 2000 GB/7.2K SATA E-DDM

RAID group types

- ▶ The DS3950 supports RAID 0, 1, 3, 5, 6, and 10.
 - Raid 0, 1, and 10 can have a maximum of 112 drives in the RAID group (array).
 - RAID 3, 5, and 6 can have a maximum of 30 drives in the RAID group (array).

DS3950 Storage Subsystem front view

Figure 2 shows the DS3950 from the front side.




Figure 2 DS3950 storage subsystem front view

The hot-swap features of the DS3950 enable you to remove and replace the hard disk drives without turning off the storage expansion enclosure. You can maintain the availability of your system while a hot-swap device is removed, installed, or replaced.

The DS3950 Express supports up to 16 disk drive modules (DDMs). Each drive bay also provides dual paths to each drive for path redundancy. The drives are customer replacement units (CRUs).

Several LED indicators and the FC Link speed selector are also visible from the front of the storage unit. For a detailed description of all the LEDs and their use see the DS5020 section of *IBM Midrange System Storage Hardware Guide*, SG24-7676.

Attention: Never hot-swap an Enhanced Disk Drive Module (E-DDM) CRU when its associated green activity LED is flashing.

Hot-swap a drive CRU only when its associated amber fault LED light is not flashing or when the drive is inactive and its associated green activity LED light is not flashing. Wait 70 seconds before inserting the drive back into the bay.

DS3950 storage subsystem rear view

The rear of the DS3950 can be ordered in two different versions, depending on which host port configuration is needed for the solution.

Model 94H, shown in Figure 3, does not have the iSCSI Host Interface Cards (HIC) and it comes with four 8 Gbps Fibre Channel host ports.

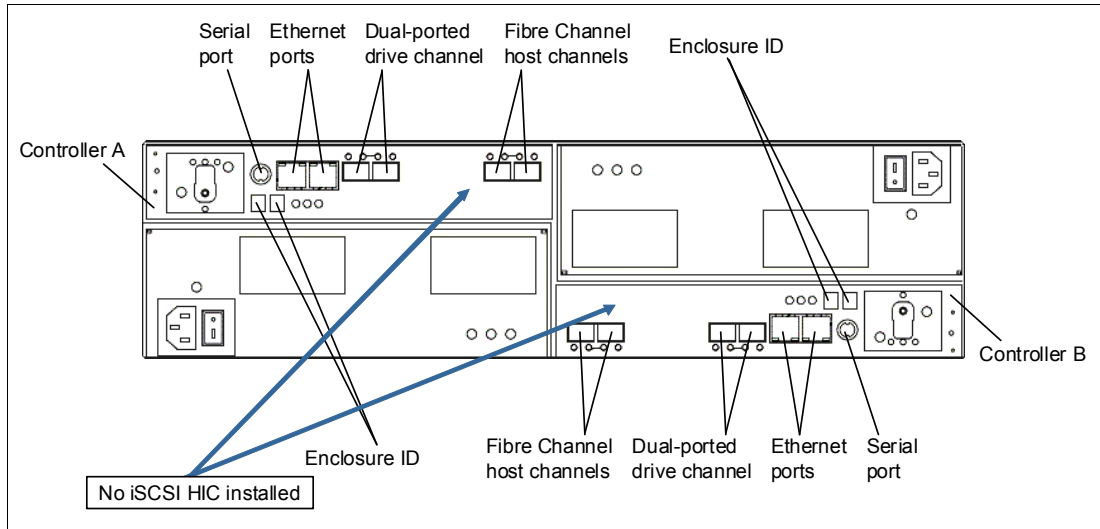


Figure 3 DS3950 rear view: four FC ports (model 94H)

Model 98H, shown in Figure 4, has an additional four 1 GBps iSCSI host ports along with the four 8 Gbps Fibre Channel host ports (Figure 4).




Figure 4 DS3950 rear view: four FC and four iSCSI ports (model 98H)

The DS3950 rear view shows four hot-swappable parts:

- ▶ The two controllers with the Backup Battery Unit (BBU)
- ▶ The two Power Supply and Fan Units

The two controllers hold host and drive interfaces as well as the batteries. The left controller is controller A and the right controller is controller B.

Note that controller A is upside-down relative to controller B. The same configuration applies to the power supply and fan unit. It is important to keep this information in mind when connecting the back-end ports to hosts and drive-side expansion enclosures.

DS3950 components

Figure 5 shows the component breakdown of the DS3950 storage subsystem.




Figure 5 DS3950 hardware components

AC power supply and fan unit

As shown in Figure 6, there are four LEDs on the AC power supply and fan unit. These indicators provide the following status:

- ▶ Green LEDs indicate a normal operating status.
- ▶ Amber LEDs (Needs Attention) indicate a possible failure.
- ▶ A blue LEDs on a CRU indicates that it is safe to remove the component.

In normal operation, only the green LEDs will be on.

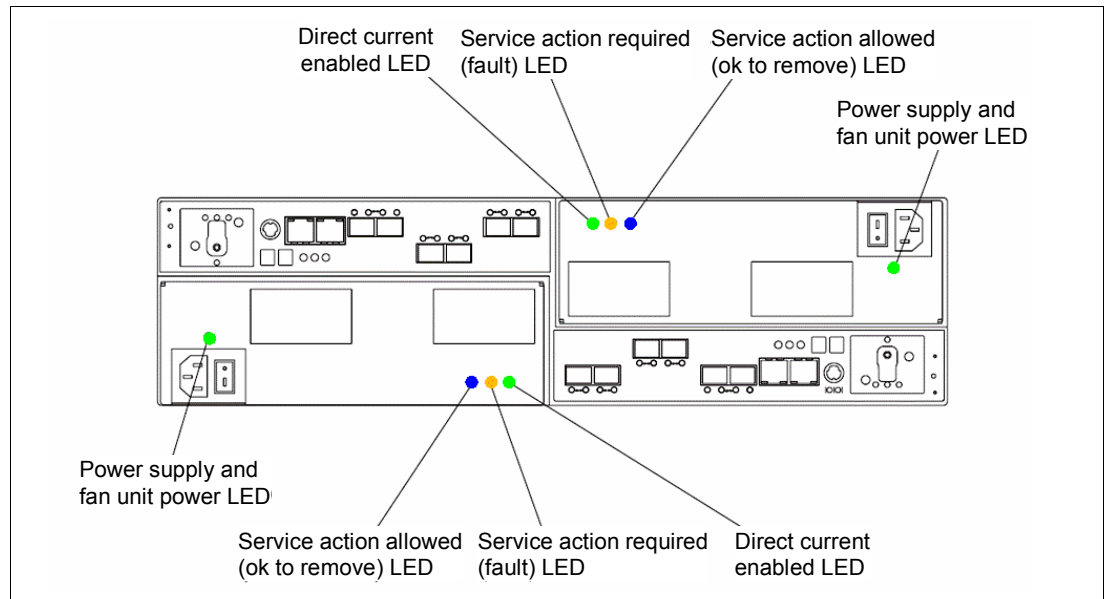


Figure 6 DS3950 power supply and fan unit LEDs

Controller module

The controllers contain the storage subsystem control logic, interface ports, and LEDs.

Depending on which DS3950 configuration you purchased, your controllers are one of the following types:

- ▶ Controllers with 1 GB memory and two standard 8 Gbps FC host ports
- ▶ Controllers with 2 GB memory, two 8 Gbps FC host ports, and one 2-port 1 GB iSCSI host card

The base controller module with the optional iSCSI HIC included is shown in Figure 7. The optional HIC card will be factory installed when the order is shipped. There are no field upgrades to allow the addition of the HIC cards.

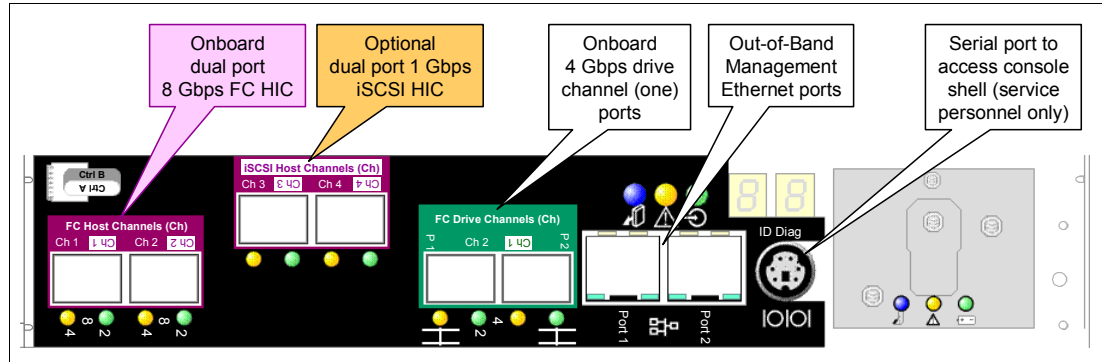


Figure 7 DS3950 controller with iSCSI host ports

Enhanced Disk Drive Modules (E-DDMs)

The hot-swap drive bays that are accessible from the front of your storage subsystem are shown in Figure 8.




Figure 8 DS3950 hot-swap drive bays

The DS3950 supports both Fibre Channel (FC) and SATA E-DDMs intermixed in the storage subsystem drive chassis. The DS3950 supports up to sixteen 4 Gbps FC or 3 Gbps SATA E-DDMs.

SATA E-DDMs have an ATA translator card (interposer card) that converts the Fibre Channel protocol interface of the DS3950 drive channel or loop into the hard drive SATA protocol interface. It also provides dual paths to the SATA drive for drive CRU path redundancy. Each drive, ATA translator card, and carrier assembly is called SATA E-DDM CRUs. The Fibre Channel E-DDMs consist of the Fibre Channel and the carrier assembly (drive tray).

Install E-DDM CRUs in the 16 drive bays on the front of the storage subsystem from the leftmost slot (slot 1) to the rightmost slot (slot 16). When an E-DDM is installed, the drive and tray slot designation is set automatically. The hardware addresses are based on the enclosure ID, which is set by the controller software, and on the E-DDM physical location in the storage subsystem.

The DS3950 storage subsystem drive channel operates at a 4 Gbps Fibre Channel interface speed. Even the 3 Gbps SATA E-DDMs operate at 4 Gbps Fibre Channel speed.

Link rate speed: Even though the DS3950 has a 2 or 4 Gbps Fibre Channel Link Rate switch that can be used to set the drive channel speed at 2 Gbps, the link rate speed *must* be set to 4 Gbps. The DS3950 supports only 4 Gbps FC speed in the drive channel.

The Link Rate switch on the DS3950 storage subsystem and all storage expansion enclosures connected to it must have the same setting.

The IBM System Storage DS3950 storage subsystem (Machine Type 1814-20A) supports RAID levels 0, 1, 3, 5, and 6 up to over 67.2 TB when using 600 GB Fibre Channel hard drives, and up to 224 TB when using 2 TB Serial Advanced Technology Attachment (SATA) Enhanced Disk Drive Modules (E-DDMs).

The DS3950 supports configurations of FC disks, SATA disks, or a mix of FC and SATA disk drives. FC disks are allocated to applications that demand high performance and have high I/O rates, and less-expensive SATA disks are allocated to applications that require less performance.

Backup Battery Unit (BBU)

The Backup Battery Unit provides power to back up the cache memory of each controller onto flash drives in the event of a power failure. Each battery unit contains a sealed, rechargeable SMART lithium ion battery. The battery unit contains enough charge to back up the cached data in each controller to a flash drive in the event of a power failure.

DS3950 storage subsystem LED indicator lights

LED indicator lights allow the DS3950 to communicate with the user. There are four main components with LEDs:

- ▶ Front panel
- ▶ RAID controllers
- ▶ Battery
- ▶ Power supply fans

For detailed description of all the LEDs and their use, see the DS5020 section of *IBM Midrange System Storage Hardware Guide*, SG24-7676.

DS3950 storage subsystem host-side connections

The DS3950 integrates the host-side and drive-side connections into the controller itself. The DS3950 has two 8 Gbps host connections by default. They are mounted on the mainboard of the controller. You can also order the optional 1 Gbps iSCSI HIC ports. See Figure 9 on page 8 for the location of all the connections shown with the optional HIC.




Figure 9 DS3950: All connectors

The FC Host connections support Fibre Channel attachment through SAN switches or direct connections. The iSCSI Host connections support 100 Mbps or 1 Gbps switched Ethernet or iSCSI network and direct connection.

- ▶ 8 Gbps FC host ports: These ports auto-negotiate with 2 Gbps, 4 Gbps, and 8 Gbps Fibre Channel speed if an 8 Gbps SFP is installed. 1 Gbps FC speed is not supported.
- ▶ 1 Gbps iSCSI host Ports: The iSCSI ports support both IPv4 and IPv6 TCP/IP addresses, CHAP, and iSNS. Use either Cat5E or Cat6 Ethernet cable types for iSCSI port connections. A Cat6 Ethernet cable provides optimal performance. The setup of the host ports will be performed in the Storage Manager. By default, the iSCSI ports auto-negotiate between 100 and 1000 Mbps Ethernet speed.

It is good practice to match host network connections to the DS3950 by attaching connections from each host network to each controller. By doing so, you benefit from the DS3950's ability to fail over and distribute the workload between the two controllers. For any given host, it is good practice to connect to the same host port number on each controller. Figure 10 on page 9 shows an example of a mixed host networking environment.

Best practice: For any host, it is best to connect to the same host port number on each controller from the host adapters. This provides ease of management and tracking.

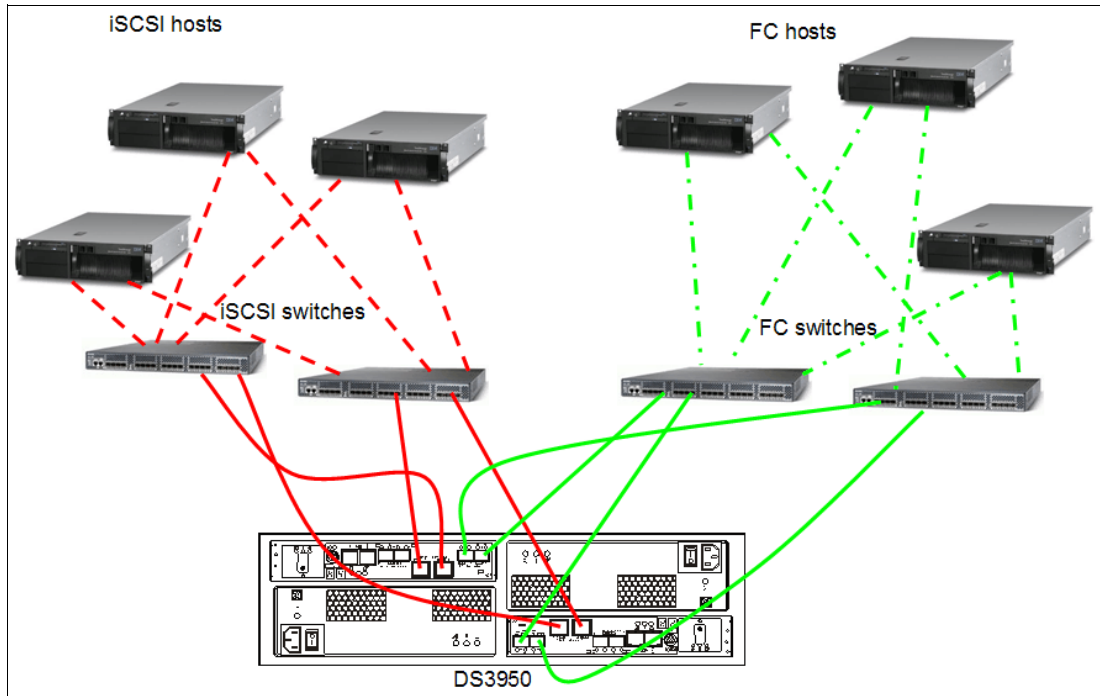


Figure 10 DS3950 mixed host connection layout

DS3950 storage subsystem drive-side connections

The DS3950 can attach up to six expansion enclosures. The regular expansion enclosure is the EXP395. However, you can buy a feature that enables you to attach EXP810 as well. Only these two expansion enclosures are currently supported.

These expansions are connected to the controller through the four (two on each controller) 4 Gbps FC drive ports. Both ports must run at 4 Gbps speed because only EXPs and drives running a 4 Gbps are supported to be attached to the DS3950. It is generally best to spread the enclosures evenly between the two drive channel pairs as you scale up the DS3950 in storage capacity. A fully configured DS3950 should have three expansion enclosures on each drive-side channel pair.

Expansion cabling rules: There are three rules for expansion cabling:

- ▶ With the DS3950, only connect a maximum of three enclosures per controller drive port.
- ▶ The DS3950 controller drive port must only be connected to the EXP395 or EXP810 port labelled 1B. Because the left (ESM A) and right (ESM B) enclosure service modules (ESM) are inserted in different orientations, ensure that you use the port labeled 1B before making the Fibre Channel connection to the DS3950 storage subsystem, as shown in Figure 11.
- ▶ Spread expansion enclosures between the two drive channel pairs. For example, if you attach four enclosures, it is better to have two enclosures behind each drive port rather than three and one.

DS3950 storage subsystem drive-side cabling

The drive-side cabling for the DS3950 depends on how many expansion units you must attach:

- ▶ If you attach only one enclosure, make sure that you have one connection to each of the controllers, thus using one of the two ports on each controller (controller A port 2 and controller B port 1), as shown in Figure 11.




Figure 11 DS3950 drive cabling with one EXP

- ▶ If you attach a second expansion unit, connect it by using the second port on the controller (controller A port 1 and controller B port 2), as shown in Figure 12.




Figure 12 DS3950 drive cabling with two EXPs

- Beyond one enclosure (up to a maximum of six), make sure that you equally distribute the enclosures among the redundant drive channel pairs by connecting the cascaded enclosure behind the first enclosure, spreading them between both channel pairs as seen in Figure 13 and Figure 14 on page 13.

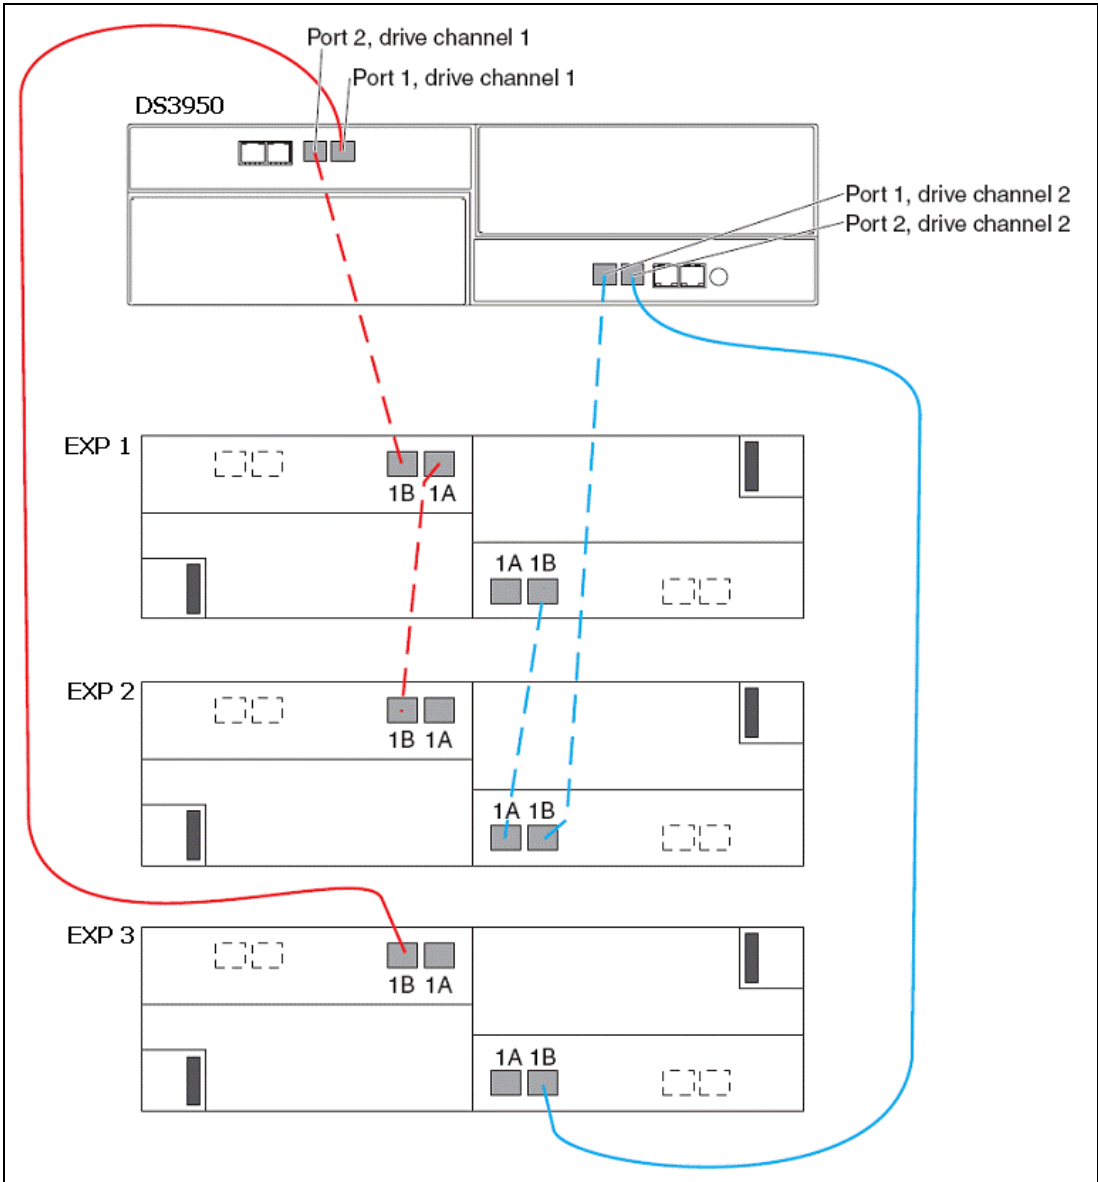


Figure 13 DS3950 drive cabling with three EXPs

When six enclosures are required, the same method is employed again, maintaining drive channel redundancy and using both controllers, as shown in Figure 14.

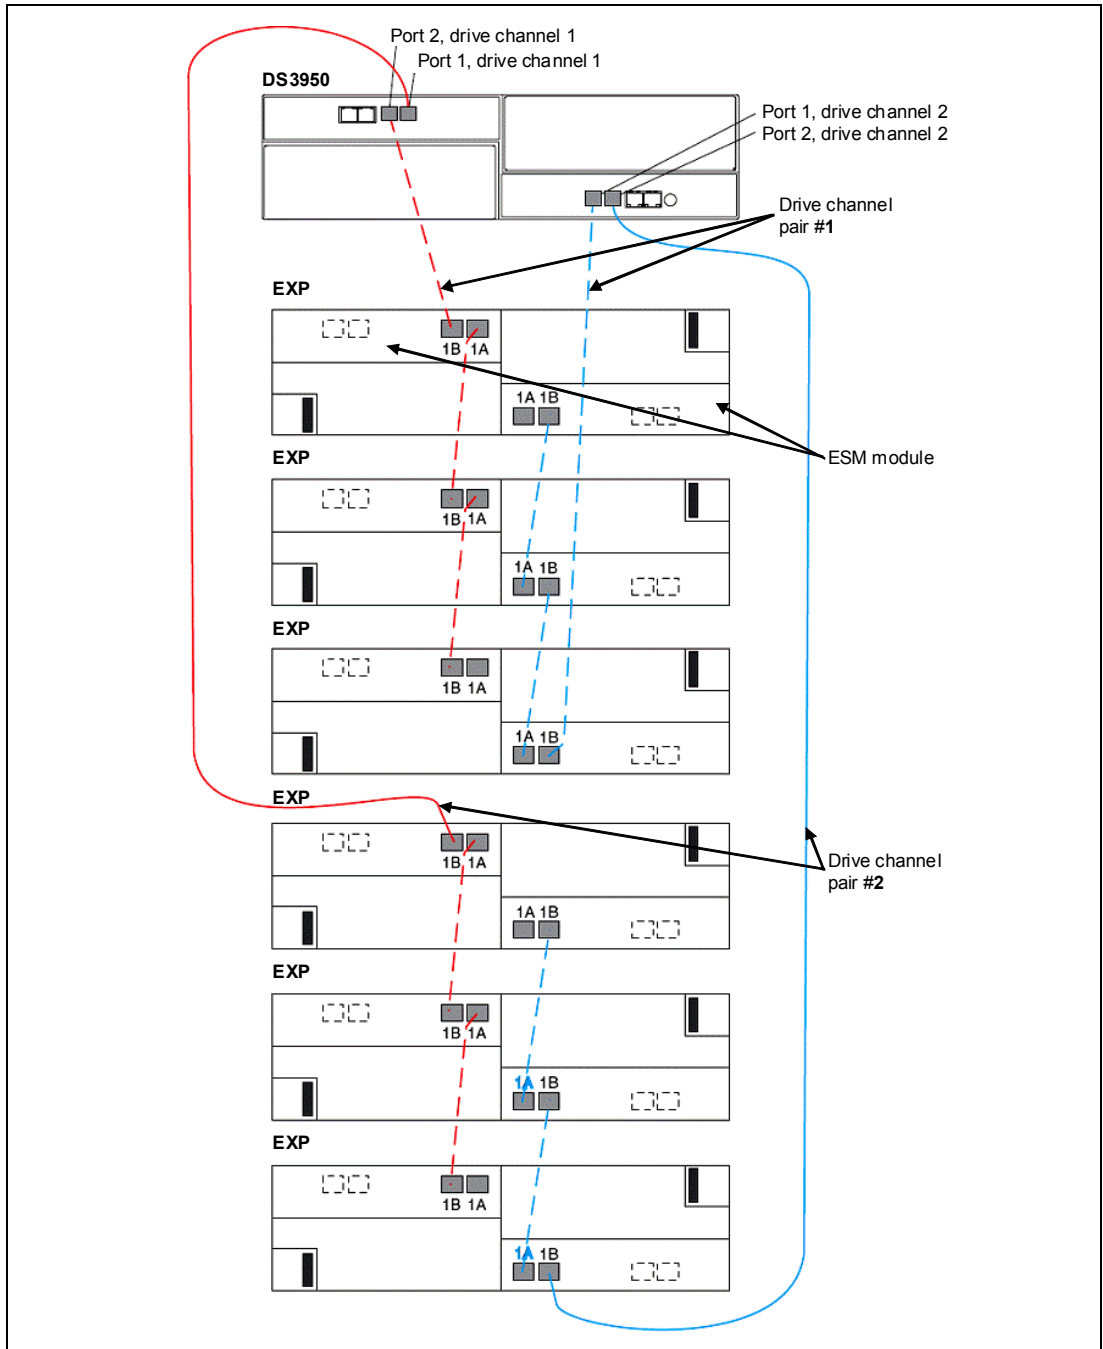


Figure 14 DS3950 drive-side cabling with six EXPs

DS3950 storage subsystem additional connections

The DS3950 storage subsystem has various kinds of connectors, as shown in Figure 3 on page 4 and Figure 4 on page 4.

These connectors are:

► **Management Ethernet connectors**

This connector is for an RJ-45 10/100/1000 BASE-Tx Ethernet connection. There are two connections per controller. One port is designed for out-of-band management and the other port is meant for serviceability.

The logic behind adding an extra port was to introduce additional isolation and to separate management and service traffic from one another. Because of the extra port, it is preferable to have two IP addresses per controller to manage and service the DS3950 storage subsystem appropriately.

However, you cannot attach this port to a routed network, because you cannot set up a gateway for it. You will still operate the DS3950 with only one IP port active per controller. Best practice is to set port 1 in the customer network for out-of-band management and keep port 2 as the default to let service personnel connect using the default IP addresses.

The default IP addresses for the controllers are listed in Table 1. The default subnet mask for all four Ethernet ports is 255.255.255.0.

Table 1 Default IP addresses for management ports

	Controller A	Controller B
Port 1	192.168.128.101	192.168.128.102
Port 2	192.168.129.101	192.168.129.102

► **Serial port**

This serial port is used for management and diagnostic purposes. You can use a PC with a terminal emulation utility, such as Hyper Terminal, to access the command set.

The maximum baud rate is 115,200 bps. The default baud rate setting from the factory is 38,400 bps, N-8-1, with no flow control.

Attention: Managing the DS3950 storage subsystem through the serial interface has potential risks. Using certain commands you can initialize the RAID controller and thereby lose all your data. Use this interface only when instructed to do so by IBM Support.

The team who wrote this paper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Sangam Racherla is an IT Specialist and Project Leader working at the International Technical Support Organization (ITSO), San Jose Center. He holds a degree in Electronics and Communication Engineering and has ten years of experience in the IT field. Sangam has been with the ITSO for the past seven years and has extensive experience installing and supporting the ITSO lab equipment for various Redbooks publication projects. His areas of expertise include Microsoft® Windows®, Linux®, AIX®, System x®, and System p® servers, and various SAN and storage products.

Alexander Watson is an ATS Specialist for Storage Advanced Technical Skills (ATS) Americas in the United States. He is a Subject Matter Expert on SAN switches and the IBM Midrange system storage products. He has more than fifteen years of experience in planning,

managing, designing, implementing, problem analysis, and tuning SAN environments and storage systems. Alexander has worked at IBM for eleven years. His areas of expertise include SAN fabric networking, Open System Storage IO, and IBM Midrange Storage solutions.

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.

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

ibm.com/redbooks/residencies.html

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>

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 give 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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

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.

This document REDP-4702-00 was created or updated on December 24, 2010.

Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:
ibm.com/redbooks
- ▶ Send your comments in an email to:
redbooks@us.ibm.com
- ▶ Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.



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:

AIX®

IBM®

Redbooks®

Redpaper™

Redbooks (logo) ®

System p®

System Storage®

System x®

The following terms are trademarks of other companies:

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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