

IBM System z Qualified DWDM Fibernet DUSAC Rel 2.2b

Bill White

Pasquale Catalano

Andrew Crimmins



z Systems



This IBM® Redpaper™ publication is one in a series that describes IBM z™ Systems qualified Dense Wavelength Division Multiplexing (DWDM) vendor products for IBM Geographically Dispersed Parallel Sysplex™ (IBM GDPS®) solutions with Server Time Protocol (STP). The protocols described in this paper are used for IBM supported solutions that require cross-site connectivity of a multisite Parallel Sysplex or remote copy technologies, which can include GDPS and non-GDPS applications. GDPS qualification testing is conducted at the IBM Vendor Solutions Connectivity (VSC) Lab in Poughkeepsie, NY.

IBM and FiberNet successfully completed qualification testing of the DUSAC 4800 platform. This paper describes the applicable environments, protocols, and topologies that are qualified and supported by IBM z Systems™ for connecting through the DUSAC 4800 platform hardware and software, Release 2.2b.

This paper is intended for anyone who wants to learn more about the DUSAC 4800 Platform Release 2.2b. This document is not meant to determine qualified products. To ensure that the planned products to be implemented are qualified, registered users can see the IBM Resource Link® library for current information about qualified DWDM vendor products:

<https://www.ibm.com/servers/resourceLink/lib03020.nsf/pages/systemzQualifiedWdmProductsForGdpsSolutions?OpenDocument&pathID=>

For more information about IBM Redbooks® publications for z Systems qualified DWDM vendor products, see this website:

<http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=qualified+AND+wdm&SearchOrder=1&SearchFuzzy=>

z Systems GDPS qualification overview

GDPS is an enterprise-wide continuous availability (CA) and disaster recovery (DR) automation solution that can manage recovery from planned and unplanned outages across distributed servers and z Systems platforms. GDPS can be configured in either a single site or in a multisite configuration. It is designed to manage remote copy configuration between storage subsystems, automate Parallel Sysplex operational tasks, and affect failure recovery. This configuration is done from a single point-of-control, which leads to improved application availability. Historically, this solution was known as a GDPS. Today, GDPS continues to be applied as a general term for a suite of business continuity solutions. This term includes solutions that do not require a dispersed or multisite sysplex environment.

GDPS supports the following forms of remote copy in multisite solutions:

- ▶ IBM System Storage® Metro Mirror, a synchronous form of remote copy, previously known as *Peer-to-Peer Remote Copy (PPRC)*
- ▶ IBM System Storage Global Mirror, an asynchronous form of remote copy for z Systems and distributed systems
- ▶ IBM System Storage z/OS® Global Mirror, an asynchronous form of remote copy for z Systems, previously known as *Extended Remote Copy (XRC)*

Depending on the form of the remote copy that is implemented, the GDPS solution is known as one of these configurations:

- ▶ GDPS with Metro Mirror
- ▶ GDPS with Global Mirror
- ▶ GDPS with z/OS Global Mirror

GDPS also offers two solutions that combine the technologies of Metro Mirror with either Global Mirror or z/OS Global Mirror. This combination allows clients to meet requirements for CA with zero data loss locally within metropolitan distances (for most failures). The combination also provides a disaster recovery solution in the case of a region-wide disaster.

The disaster recovery solutions include the following options:

- ▶ GDPS with Metro and Global Mirror (GDPS and MGM), which is a cascading data replication solution for both z Systems and distributed systems data
- ▶ GDPS with Metro and z/OS Global Mirror (GDPS with MzGM), which is a multitarget data replication solution for z Systems data

The GDPS solution is also independent of the disk vendor, if the vendor meets the specific levels of Metro Mirror, Global Mirror, and z/OS Global Mirror architectures.

For more information about GDPS, see this website:

<http://www.ibm.com/systems/z/advantages/gdps/index.html>

IBM supports DWDM products that are qualified by z Systems for use in GDPS solutions. To obtain this qualification, DWDM vendors obtain licensed IBM patents, intellectual property, and know-how that are related to the GDPS architecture. This licensing provides vendors access to the proprietary IBM protocols and applications that are used in a GDPS environment. These proprietary items include IBM Sysplex Timer, InterSystem Channel (ISC), STP, Metro Mirror, Global Mirror, and z/OS Global Mirror.

Licensing of IBM patents also provides the DWDM vendor with technical information that pertains to future IBM releases. Qualified vendors typically license this information for an extended period. This license allows them to subscribe to the latest GDPS architecture changes and to be among the first to the market with offerings that support these features.

Licensing: Check with your DWDM vendor for current IBM technology and patent licensing status.

In addition, IBM tested and qualified these vendor products with the same test environment and procedures that were used to test the protocols that provide the required connectivity of a GDPS configuration. This testing included functionality and recovery verification. Having access to these test facilities allows IBM to configure a fully functional sysplex. By using the test facilities, you can simulate failure and recovery actions that cannot be tested as part of a working client environment.

IBM has the facilities to test and qualify these products with both current and previous generation equipment within the VSC Lab. With this qualification testing, IBM can reproduce any concerns that might arise when using this equipment in a client application.

Figure 1 shows the test environment that is used for DWDM vendor qualification and the logical connections for IBM System Storage Metro Mirror and Global Mirror.

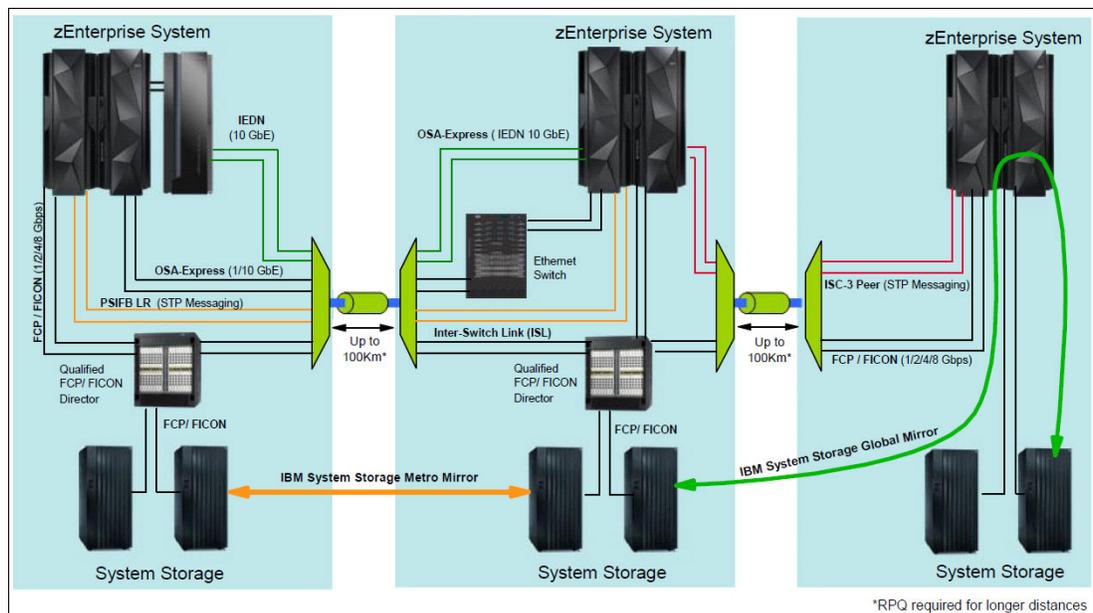


Figure 1 VSC Lab environment

Qualification testing

The VSC Lab contains z Systems hardware with software applications that test the interoperability of DWDM products within a GDPS. A typical qualification test cycle is six to eight weeks in length.

The hardware that was used for testing purposes includes (but is not limited to) the following components:

- ▶ IBM z Systems™ platforms enabled for STP
- ▶ IBM 9037 Model 2 Sysplex Timers

- ▶ IBM System Storage
- ▶ IBM z Systems qualified Fibre Connection (FICON®) Directors
- ▶ IBM Ethernet products
- ▶ IBM zEnterprise® BladeCenter Extension (10 Gigabit Ethernet)

DWDM links of varying distances are deployed by using spools of single-mode fiber in lengths 5 - 50 km (3.11 - 31.07 miles). Multiple spools are interconnected to test DWDM link protocols up to the maximum supported distances. To achieve the maximum distances that are qualified for GDPS protocols, vendors can use optical amplifiers (OAs) and dispersion compensation units (DCUs). They can insert the OAs and DCUs, at various link points, to condition the signals on the fiber links. These links are connected to the DWDM equipment.

Qualified DCUs: Fiber-based DCUs are not qualified for use in STP applications unless stated in the qualification letter.

The operating system and application software are installed to create and to stress test the GDPS environment. The software used in the test environment includes, but is not limited to, the following components:

- ▶ z/OS, Linux on z Systems, and Parallel Sysplex software exploiters
- ▶ Coupling Facility Control Code (CFCC)
- ▶ IBM proprietary software and microcode utility test suites

As part of the GDPS qualification test, IBM proprietary software and microcode utility test suites are used. The software and test suites drive the various GDPS components and protocols to the full data rate of each link type that is transported by the DWDM equipment. This level of testing ensures that the maximum channel utilization is achieved and tested to levels well beyond typical client environments.

The test suites are used for verification of z Systems architecture functionality. For a trial to be classified as successful during these functionality tests, no errors can be detected by the attached subsystems. Any errors that are detected during this testing are captured and analyzed by the test suites.

The test suites are also used for verification of z Systems architecture recovery by creating various fault and error conditions. The recovery tests check for the correct detection of a fault or error condition by the attached subsystems, and ensure that the recovery adheres to z Systems architecture rules.

Some of the recovery tests conducted for each link type include the following actions:

- ▶ Link state change interrupt detection and recovery: Links are deliberately broken and reestablished to ensure that detection and error recovery occur correctly.
- ▶ Link error threshold and link synchronization error detection and recovery: Errors are deliberately injected, at the application and channel subsystem levels, into the link protocol data streams to ensure that detection and error recovery take place correctly.
- ▶ Link service and maintenance package recovery: Link hardware maintenance actions are conducted to ensure that link state change detection and recovery take place correctly.
- ▶ Link protection schemes and recovery: Vendor-specific protection methods are tested to ensure that the expected link errors are detected and that recovery takes place correctly.
- ▶ STP timing link recovery: STP timing links are broken to ensure that changes to the Coordinated Timing Network (CTN) behave correctly when the break occurs and when the links are reconnected.

GDPS components and protocols

The IBM technologies featured in this section are functional components of GDPS and are tested during the qualification process. Clients can also use these components in environments that do not require a full GDPS solution. The testing provides a level of assurance that the components function when used with a qualified DWDM platform.

Components

The following GDPS components are tested during the qualification process:

- ▶ IBM z Systems platforms
- ▶ IBM Parallel Sysplex
- ▶ IBM System Storage
- ▶ IBM System Storage Metro Mirror (PPRC)
- ▶ IBM System Storage Global Mirror
- ▶ IBM System Storage z/OS Global Mirror (XRC)
- ▶ Vendor DWDM platform
- ▶ IBM Ethernet products
- ▶ IBM zEnterprise BladeCenter Extension
- ▶ Inter-Switch Links (ISLs) between two z Systems qualified FICON/Fibre Channel Directors

Protocols

Table 1 lists the GDPS connectivity protocols and their data transfer rates. Not all protocols were tested on the DUSAC 4800 platform. For a complete list of the protocols and interface cards that were qualified, see Table 2 on page 15.

Table 1 GDPS supported protocols

Protocol	Data transfer rate
Enterprise Systems Connection (ESCON) ^a	200 Mbps ^b
Control Link Oscillator (CLO) ^a	8 Mbps
External Time Reference (ETR) ^a	8 Mbps
Fibre Connection (FICON)	1 Gbps
Fibre Connection (FICON) Express2	1, 2 Gbps
Fibre Connection (FICON) Express4	1, 2, 4 Gbps
Fibre Connection (FICON) Express8	2, 4, 8 Gbps
Fibre Channel FC100/ FC200/ FC400/ FC800	1, 2, 4, 8 Gbps
Inter-Switch Link (ISL) FC100/ FC200/ FC400/ FC800/ FC1000/ FC1600	1, 2, 4, 8, 10, 16 Gbps
InterSystem Channel-3 (ISC-3) Peer Mode ^c	2 Gbps
STP (ISC-3 Peer Mode with STP message passing) ^c	2 Gbps
STP (Parallel Sysplex InfiniBand Long Reach (PSIFB LR) 1x IB-single data rate (SDR) with STP message passing)	2.5 Gbps

Protocol	Data transfer rate
STP (PSIFB LR 1x IB-double data rate (DDR) with STP message passing)	5 Gbps
Gigabit Ethernet (GbE)	1, 10 Gbps

- a. ETR, CLO, and ESCON are not supported on the zEC12 and zBC12.
- b. Effective channel data rate of an ESCON channel is affected by distance.
- c. The zEC12 and zBC12 are the last platforms to support InterSystem Channel-3 (ISC-3).

Often, these supported protocols are also used in environments that are not GDPS. Robust testing is conducted during the qualification process. The results must provide clients with a high level of confidence when using these z Systems qualified DWDM vendor platforms in environments that are not GDPS.

Server Time Protocol

Server Time Protocol (STP) is designed to provide the capability for multiple servers and Coupling Facilities (CFs) to maintain time synchronization with each other without requiring an IBM Sysplex Timer. STP can help provide functional and economic benefits when compared to the Sysplex Timer. STP allows concurrent migration from an ETR network and can coexist with an ETR network.

STP is a message-based protocol in which STP timekeeping information is passed over externally defined coupling links: InterSystem Channel-3 (ISC-3) Peer Mode and PSIFB links. ISC-3 links in peer mode and PSIFB Long Reach (LR) are the only coupling links that can be used to transport STP messages between data centers over a DWDM platform.

The STP design introduced a concept called Coordinated Timing Network (CTN), which is a collection of servers and CFs that are time synchronized to a time value. This value is called Coordinated Server Time (CST). The CST represents the time for the entire network of servers. A CTN can be configured as either an STP-only CTN or a Mixed CTN. For more information, see *Server Time Protocol Planning Guide*, SG24-7280, *Server Time Protocol Implementation Guide*, SG24-7281, and *Server Time Protocol Recovery Guide*, SG24-7380.

Figure 2 shows a multisite STP-only CTN.

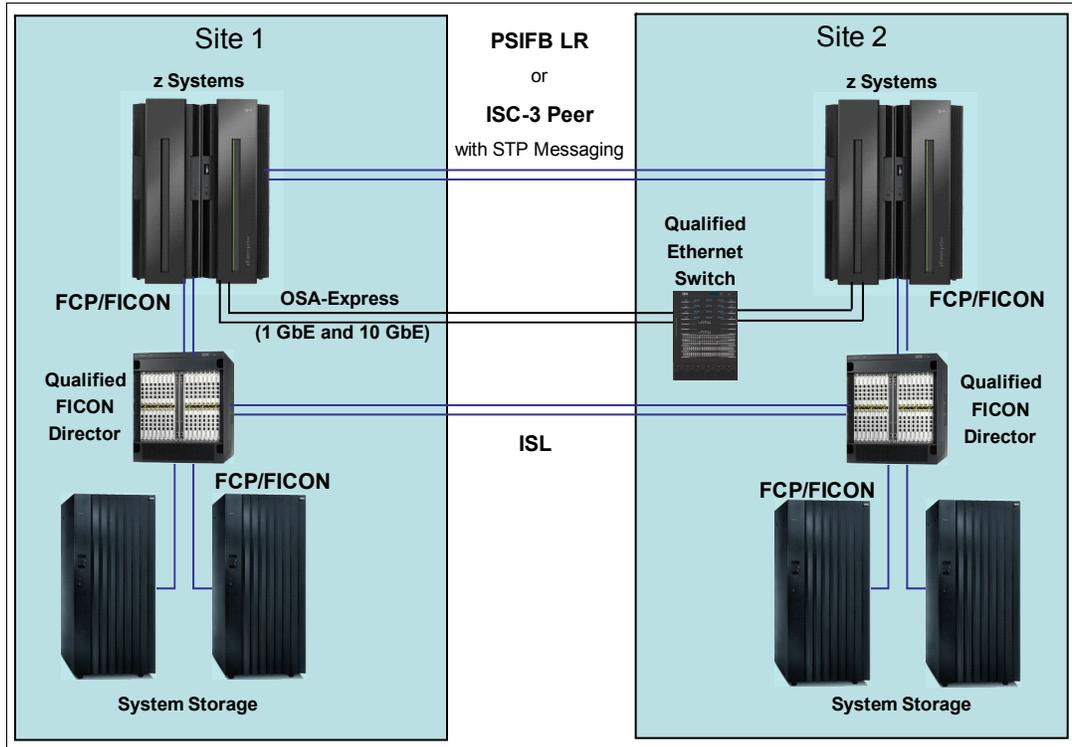


Figure 2 Multisite STP-only CTN

IBM Sysplex Timer

The Sysplex Timer is a mandatory component of an ETR network and an STP Mixed CTN. The Sysplex Timer provides an ETR to synchronize the time-of-day (TOD) clocks on the attached z Systems platforms in a GDPS or PPRC environment.

Figure 3 shows a multisite ETR network.

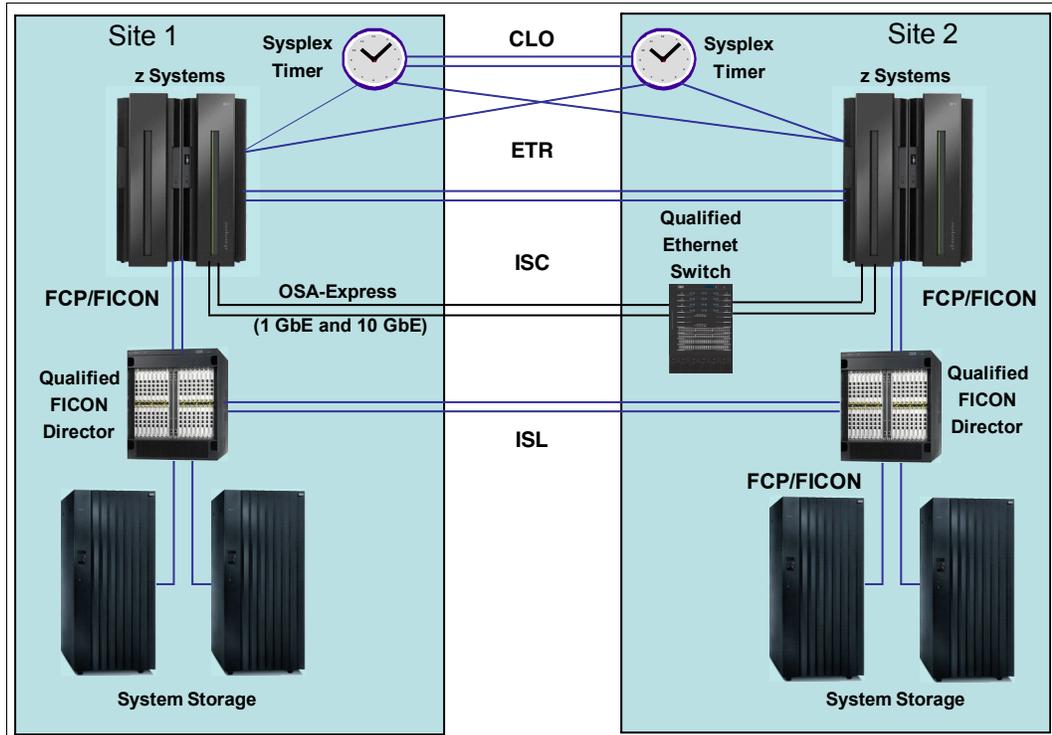


Figure 3 Multisite ETR network with Sysplex Timers

The Sysplex Timer uses two link types:

- ▶ External Time Reference (ETR)
 - ETR links are connections between the Sysplex Timer and the z Systems ETR ports that provide TOD clock synchronization between multiple servers.
- ▶ Control Link Oscillator (CLO)
 - CLO links are connections between two Sysplex Timer units in an expanded availability configuration that allow synchronization of the Sysplex Timer timing signals.

IBM zEnterprise BladeCenter Extension

The IBM zEnterprise BladeCenter Extension (zBX) is available as an option with the IBM zEnterprise System. The zBX brings the computing capacity of systems in blade form-factor to the zEnterprise System and provides several distributed environments, such as IBM AIX® on POWER7®, Linux on System x, and Microsoft Windows on System x.

The zBX is designed with a redundant hardware infrastructure that supports the multi-platform environment in a seamless, integrated way. Its hardware features are part of the mainframe, not add-ons.

The zBX and the zEnterprise central processor complex (CPC) are interconnected via a high-speed private network called an intraensemble data network (IEDN), which consists of top-of-rack 10 GbE switches in the zBX and OSA-Express 10 GbE features in the zEnterprise CPC.

The IEDN provides private and secure data paths between all elements of a zEnterprise ensemble. A zEnterprise ensemble is a collection of highly virtualized diverse systems that can be managed as a single logical entity, and where diverse workloads can be deployed. It is composed of up to eight nodes, each of which consists of a zEnterprise CPC and an optional zBX (see Figure 4).

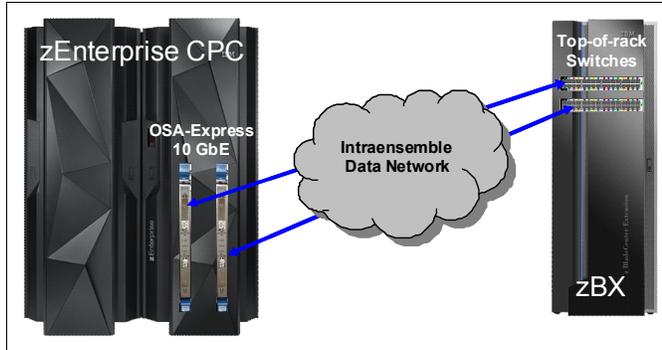


Figure 4 zEnterprise node

GDPS/PPRC and GDPS/GM support zBX hardware components, providing workload failover for automated multi-site recovery. These capabilities can help facilitate the management of planned and unplanned outages across the zEnterprise System.

For more information regarding the zBX, see *Building an Ensemble Using IBM zEnterprise Unified Resource Manager*, SG24-7921.

IBM 10 Gigabit Ethernet (10GbE) RoCE Express

The 10GbE RoCE Express feature exploits Remote Direct Access Memory over Converged Ethernet (RoCE) and is designed to provide very fast memory-to-memory communications between two zEnterprise CPCs.

Use of the 10GbE RoCE Express feature helps reduce consumption of CPU resources for applications utilizing the TCP/IP stack, and might also help to reduce network latency with memory-to-memory transfers utilizing Shared Memory Communications - Remote Direct Memory Access (SMC-R).

SMC-R uses a TCP connection to determine eligibility to exploit the RoCE fabric and then to build a point-to-point SMC-R link for the data flow over the RDMA connection path. TCP/IP is used not only to establish the TCP/IP and RDMA connections, but also for Keepalive functions and to terminate the TCP and the associated RDMA connections. The standard TCP/IP path over an OSA-Express port need not be dedicated to RDMA usage; it can be used simultaneously for other, non-RDMA traffic.

The OSA-Express port and the 10GbE RoCE Express port can be attached to a 10 Gigabit Ethernet switch or they can both have a point-to-point connection (see Figure 4).

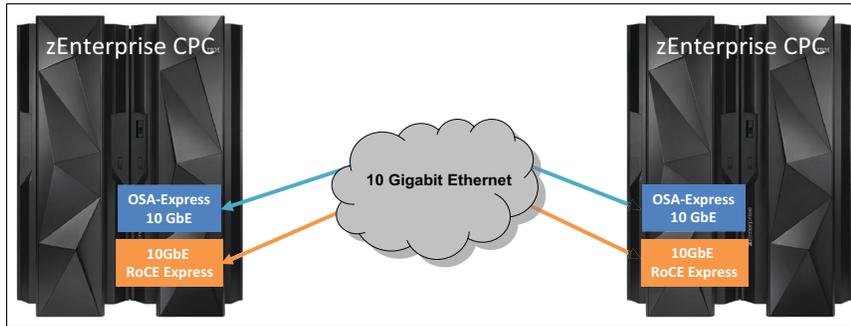


Figure 5 RoCE in a two site configuration via DWDM

The 10GbE RoCE Express feature uses a short reach (SR) laser as the optical transceiver and supports the use of a multimode fiber optic cable terminated with an LC Duplex connector.

If the IBM 10GbE RoCE Express features are connected to 10 GbE switches, the switches must support the following requirements:

- ▶ Global Pause function enabled
- ▶ Priority Flow Control (PFC) disabled
- ▶ No firewalls or routing

Note: The 10GbE RoCE Express feature is only supported on the zEC12 and zBC12.

For more information about the RoCE feature, see *IBM z Systems Connectivity Handbook*, SG24-5444.

Connectivity considerations with DWDM

When planning DWDM intersite connectivity for GDPS environments, consider the following items:

- ▶ Differential delay: The difference in the distance or latency between a transmit fiber and a receive fiber in a single fiber pair of a timing link. When configuring, adhere to the following requirements:
 - To ensure correct server time synchronization, examine the end-to-end lengths of the transmit and receive fibers within an individual STP link (ISC-3 Peer Mode with STP messaging or PSIFB). The lengths must not exceed the equivalent of *900 meters* differential delay between the transmit and receive paths of the link. This stipulation includes all dense wavelength division multiplexing (DWDM) components, OAs, DCUs¹, dark fiber links, and any time-division multiplexing (TDM)-based aggregation.

¹ Fiber-based DCUs are not qualified for use in STP applications unless stated in the qualification letter.

- To ensure correct Sysplex Timer and server time synchronization, the end-to-end lengths of the transmit and receive fibers within an individual ETR or CLO link must be equal (within *10 meters*). This stipulation includes all DWDM components, OAs, DCUs, and dark fiber links.

Supported devices: DCUs might contain significant lengths of fiber that must be included in the differential delay calculation for timing links. Not all DCUs are supported. We suggest that you check the qualification letters for supported devices.

- ▶ If both the CF data and STP timing information must be transmitted between two servers, you cannot select a subset of coupling links to be used just for STP timing information.
- ▶ Protection schemes, for individual timing links or for fiber trunks that transport timing links, must be bidirectional (switch both transmit and receive paths in the event of a failure). This configuration ensures that the correct differential delay is maintained.
- ▶ DWDM configurations must have high availability topologies in place to ensure that there are no single points of failure.
- ▶ Fiber trunk protection schemes must be designed with two trunk switching modules and four site-to-site fiber pairs that are carried over at least two diverse routes. STP, ETR, and CLO links must connect by using separate trunk switching modules to ensure that a fiber trunk protection event does not interrupt all timing links simultaneously.
- ▶ TDM-based aggregation DWDM hardware can be used for STP links (ISC-3 Peer Mode or PSIFB LR with STP message passing) only if it is qualified for STP usage.
- ▶ GDPS DWDM configuration must be a point-to-point fixed dark fiber network.

For more information about the STP recovery process, see the *Server Time Protocol Recovery Guide*, SG24-7380.

Technical description

The DUSAC 4800 Platform is a scalable, high-speed fiber-optic data transport system. It consists of a modular chassis, which can be interconnected to form an optical network that supports International Telecommunications Union (ITU)-specific wavelengths. The wavelengths are multiplexed onto a single fiber by using DWDM. The DUSAC 4800 also supports the use of a single pair of fibers multiplexing solution.

The optical network that connects the DUSAC 4800 Platform can be configured in a two-site point-to-point, multisite ring, or meshed network.

Note: IBM only qualifies point-to-point topologies for GDPS solutions. Other topologies might contain significant length variations of fiber that go beyond the differential delay limits for timing links and are therefore not supported.

The DUSAC 4800 Platform was qualified by using software Release 2.2b. It has a modular chassis that can house multiple optical interface cards, depending on the chassis. Multiple DUSAC 4800 chassis can be interconnected to support larger quantities of client interfaces.

Interface cards and modules

This section includes a detailed list of the DUSAC 4800 optical interface cards and modules that are qualified by z Systems GDPS testing.

The following optical interface cards and modules are qualified by z Systems GDPS testing:

► Transponder cards

Transponder cards connect client equipment to the DWDM platform. The transponder card converts client optical signals to an ITU-compliant DWDM wavelength for transmission to the remote site.

DUSAC 4800 Transponder cards have pluggable client optical transceivers.

The following DUSAC 4800 Transponder card types are qualified:

- FTX-n: Single port Transponder card for services up to 10 Gbps
- FTX-10C: Single port Encryption Transponder card for services up to 10 Gbps
- FTX-16S: Single port Transponder card for 16 Gbps services
- FTX-6: Single port Transponder card for services up to 2.5 Gbps

► Optical Filter Modules

The Optical Filter Modules are optical components that combine (multiplex) or separate (demultiplex) multiple optical signals onto a fiber link.

The reconfigurable optical add-drop multiplexer (ROADM) allows adding or dropping of optical signals at the wavelengths layer without an optical-electrical-optical conversion in multiple directions in an optical cross-connect system.

The DUSAC 4800 supports a series of optical filter modules that can multiplex or demultiplex signals onto a fiber pair or a single fiber for transmission between sites. Optical filter modules are used in a tiered arrangement to multiplex or demultiplex groups of client channels on a fiber link for transmission between sites.

► Optical Amplifier (OA)

The OA module provides in-line optical amplification of multiple ITU-specific wavelengths within a fiber. For particular extended distances, the use of these OAs might be required.

► Dispersion Compensation Unit (DCU)

The DCU provides compensation to chromatic dispersion fiber impairment, which causes signal degradation and limits the transmission distance.

- The DCM_4800²: DCU based on Fiber Bragg Gratings is a qualified DCU type of the DUSAC 4800 platform.

Delay calculation for timing links: OAs and DCUs contain significant lengths of fiber, which must be included in the differential delay calculation for timing links.

► Protection Switch Module (PSM)

The DUSAC 4800 supports optically switched fiber protection for point-to-point DWDM links with the use of the PSM. The PSM protects site-to-site traffic from physical damage to a fiber or fiber pair by switching bidirectionally to a redundant fiber or fiber pair.

The following FiberNet DUSAC 4800 Platform PSM types is qualified:

OSW-3: Optical Protection Switch. Bi-directional transmission within a single fiber (qualified for use with STP). The OSW-3 also incorporates features to detect degradation and intrusions of the fiber link, reports and statistics can be generated.

² the DCM_4800 is qualified for STP.

Topologies and protection schemes

z Systems qualifies a two-site point-to-point DWDM network topology and protection against failures in site-to-site fiber links or failures in individual components within the DWDM network for GDPS.

GDPS is a high-availability solution that can use several protection schemes. Some restrictions apply for particular protocols, for example, Sysplex Timer (ETR and CLO), and STP (ISC-3 Peer Mode and PSIFB LR) links.

Important: Protection schemes must guarantee the correct differential delay for individual timing links:

- ▶ STP links must not exceed the equivalent of *900 meters* differential delay between the transmit and receive paths of the link.
- ▶ ETR and CLO links must not exceed the equivalent of *10 meters* differential delay between the transmit and receive paths of the link.

Protection schemes

The DUSAC 4800 platform provides the following protection schemes:

▶ Unprotected

An unprotected transponder or Muxponder card is connected to one client interface and to one site-to-site fiber link only. A failure of the transponder or Muxponder card or the site-to-site fiber link results in a loss of client communications.

▶ Client-based protection

Client-based protection uses at least two client interfaces that are connected to the DWDM. These interfaces are arranged so that the transponder or Muxponder cards that connect the two sites are distributed over two diverse site-to-site fiber links. The client device is responsible for ensuring that a failure of a DWDM module or of a single site-to-site fiber link does not result in a total loss of client communications.

GDPS timing links (ETR, CLO, and STP) are qualified for use in a client-based protection scheme if they are using separate paths (routes).

Diagram components: For simplicity, several of the components in the optical path are not shown in the diagrams in this section.

Figure 6 shows a high-level view of the client-based protection scheme. In this case, a client device has two separate site-to-site connections: one through Card A and the other through Card B. Transponder card-based and Muxponder card-based schemes for the DUSAC 4800 are shown separately.

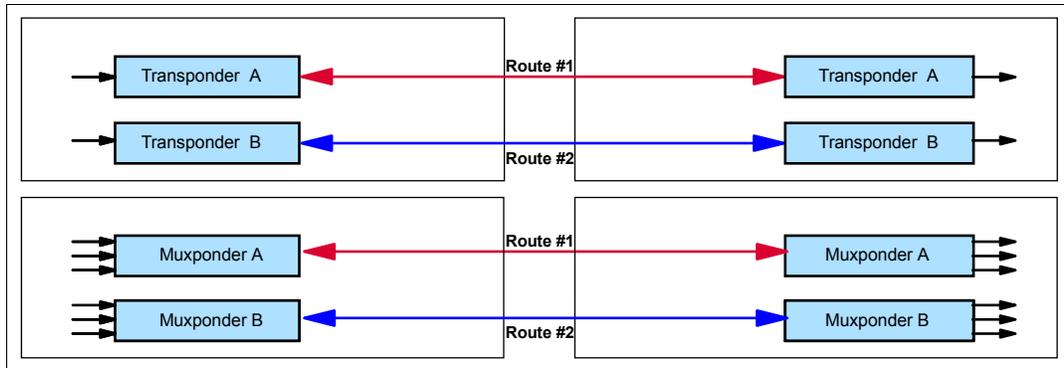


Figure 6 Client-based protection scheme

► **Optically Switched Fiber protection**

The Protection Switch Module (PSM) provides protection at the site-to-site fiber level. The PSM protects all wavelengths that are being carried on a fiber pair simultaneously. If a site-to-site fiber failure occurs, all traffic is switched bidirectionally to the backup link.

The PSM is only available for point-to-point DWDM network topologies.

Figure 7 shows a high-level view of the Optically Switched Fiber protection scheme.

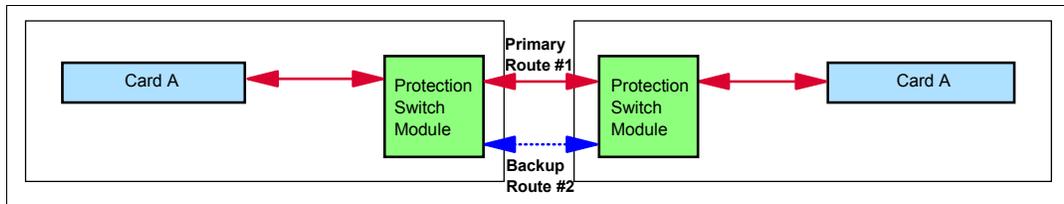


Figure 7 Optically Switched Fiber protection

GDPS timing links (STP, ETR, and CLO) are not qualified for use with a single PSM at each site. A single PSM-based, Optically Switched Fiber protection scheme must not be used with GDPS or PPRC. It must be used in conjunction with the client-based protection scheme.

If a GDPS or PPRC solution is to use Optically Switched Fiber protection, dual PSMs at each site with four site-to-site fiber pairs (trunks) are suggested. The GDPS timing links must connect by using separate cards and PSMs to ensure cross site connectivity is not lost during a switch-over.

Figure 8 shows a high-level view of the dual Optically Switched Fiber protection scheme.

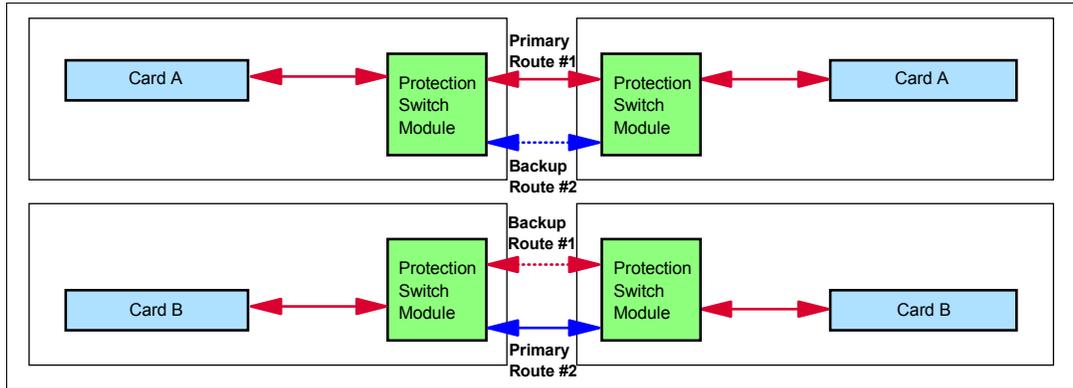


Figure 8 Dual Optically Switched Fiber protection

Protection scheme intermix

All protection schemes can be intermixed within the same DUSAC 4800 chassis or network on an individual client interface basis.

Interface card specifications

Table 2 lists the specifications of the qualified DUSAC 4800 interface cards and tested protocols.

Support note: DUSAC 4800 interface cards do not support auto-negotiation of the link speeds. For z Systems FICON and Fibre Channel Protocol (FCP) client links, the wanted link speed must be configured in the DWDM card client interface at both ends of the link.

For particular extended distances, the use of OAs and DCUs might be required. For distance and link budget specifications, reference the DWDM vendor documentation.

Table 2 Qualified client interface card details

Card type and protocol	Fiber type	Light source	Qualified distance ^{a,b}
FTX-n: Single port Transponder card for services up to 10 Gbps^c			
Fibre Channel (1, 2, 4, 8 Gbps)	SM	1310 nm	100 km (62 miles)
Fibre Channel (1, 2, 4, 8 Gbps)	MM	850 nm	100 km
Ficon (1, 2, 4, 8 Gbps)	SM	1310 nm	100 km
Ficon (1, 2, 4, 8 Gbps)	MM	850 nm	100 km
ISL (1, 2, 4, 8, 10 Gbps) ^d	SM	1310 nm	100 km
ISL (1, 2, 4, 8, 10 Gbps) ^d	MM	850 nm	100 km
ISC-3 Peer Mode (2 Gbps) ^e with STP ^f	SM	1310 nm	100 km
PSIFB 1x IFB-DDR LR (5 Gbps) ^e with STP	SM	1310 nm	100 km
Gigabit Ethernet (1, 10 Gbps) ^g	SM	1310 nm	100 km
Gigabit Ethernet (1, 10 Gbps) ^{g,h}	MM	850 nm	100 km

Card type and protocol	Fiber type	Light source	Qualified distance ^{a,b}
FTX-10C: Single port Encryption Transponder card for services up to 10 Gbps^c			
Gigabit Ethernet (10 Gbps) ^g	SM	1310 nm	100 km
Gigabit Ethernet (1, 10 Gbps) ^{g,h}	MM	850 nm	100 km
FTX-16S: Single port Transponder card for 16 Gbps services^c			
ISL (16 Gbps) ^d	SM	1310 nm	100 km
ISL (16 Gbps) ^d	MM	850 nm	100 km
FTX-6: Single port Transponder card for services up to 2.5 Gbps^c			
CLO (8 Mbps)	MM	1310 nm	40 km (24 miles)
ETR (8 Mbps)	MM	1310 nm	100 km
ESCON (200 Mbps)	MM	1310 nm	100 km
SM = single-mode fiber (9/125 micron) MM = multimode fiber (50/125 or 62.5/125 micron)			

- a. Qualified distance is based on the VSC Lab test environment, which is a physical point-to-point topology without any intermediate nodes besides OAs and DCUs.
- b. Requires request for price quotation (RPQ) - 8P2263 (IBM z9@ EC, z9 BC, IBM z10™ EC), 8P2340 (z10 BC, z196, z114), 8P2581 (zEC12), 8P2781 (zBC12) for distances over 100 km.
- c. Protocol and wavelength support depend on a pluggable client interface transceiver.
- d. Protocol is configured as FCP on DWDM.
- e. Multiple STP links can be supported on the same card.
- f. ISC-3 without STP is also supported.
- g. In addition, the 10 GbE connection is also qualified with the IBM zEnterprise intraensemble data network (IEDN).
- h. On this card, the 10 GbE connection is also qualified for RoCE (RDMA over Converged Enhanced Ethernet) using Shared Memory Communications - RDMA (SMC-R) on the zEC12 and zBC12.

Qualified distances: Consult your storage area network (SAN) switch vendors for qualified ISL-supported distances.

References

For more information about z Systems connectivity, see these resources:

- ▶ z Systems I/O connectivity home page
<http://www.ibm.com/systems/z/hardware/connectivity/index.html>
- ▶ *IBM System z Connectivity Handbook*, SG24-5444
<http://www.redbooks.ibm.com/abstracts/sg245444.html?Open>
- ▶ *FICON Planning and Implementation Guide*, SG24-6497
<http://www.redbooks.ibm.com/abstracts/sg246497.html?Open>
- ▶ *Implementing and Managing InfiniBand Coupling Links on System z*, SG24-7539
<http://www.redbooks.ibm.com/abstracts/sg247539.html?Open>

For more information about GDPS, see these resources:

- ▶ GDPS home page
<http://www.ibm.com/systems/z/advantages/gdps/index.html>
- ▶ *GDPS Family: An Introduction to Concepts and Capabilities*, SG24-6374
<http://www.redbooks.ibm.com/abstracts/sg246374.html?Open>

For more information about STP, see these resources:

- ▶ Parallel Sysplex home page:
<http://www.ibm.com/systems/z/advantages/psf/index.html>
- ▶ *Server Time Protocol Planning Guide*, SG24-7280
<http://www.redbooks.ibm.com/abstracts/sg247280.html?Open>
- ▶ *Server Time Protocol Implementation Guide*, SG24-7281
<http://www.redbooks.ibm.com/abstracts/sg247281.html?Open>
- ▶ *Server Time Protocol Recovery Guide*, SG24-7380
<http://www.redbooks.ibm.com/abstracts/sg247380.html?Open>

For more information about optical communications, see this website:

- ▶ *Understanding Optical Communications*, SG24-5230
<http://www.redbooks.ibm.com/abstracts/sg245230.html?Open>

For current information about qualified DWDM vendor products, registered users can see the library at the IBM Resourcelink website:

<https://www.ibm.com/servers/resourcelink/lib03020.nsf/pages/systemzQualifiedWdmProductsForGdpsSolutions?OpenDocument&pathID=>

For current information about qualified vendor switches and directors for IBM z Systems FICON and FCP channels, registered users can see the library at the Resourcelink website:

<https://www.ibm.com/servers/resourcelink/lib03020.nsf/pages/switchesAndDirectorsQualifiedForIbmSystemZRFiconRAndFcpChannels?OpenDocument>

For more information about IBM Redbooks publications on z Systems qualified DWDM vendor products, see this website:

<http://www.redbooks.ibm.com/cgi-bin/searchsite.cgi?query=qualified+AND+wdm&SearchOrder=1&SearchFuzzy=>

For more information about the DUSAC 4800 platform, see this website:

<http://www.fibernet.es/en/>

The team who wrote this Redpaper

This Redpaper publication was produced by a team of specialists working at the IBM International Technical Support Organization, Poughkeepsie Center.

Bill White is a Project Leader at the IBM International Technical Support Organization, Poughkeepsie, NY.

Pasquale Catalano is a Senior Engineer at the IBM Vendor Solutions Connectivity Lab in Poughkeepsie, NY. He has 11 years of experience in DWDM vendor qualification testing. He has a Masters of Science in Electrical Engineering from the State University of New York at New Paltz.

Andrew Crimmins is an Engineer at the IBM Vendor Solutions Connectivity Lab in Poughkeepsie, NY. He has four years of experience in DWDM vendor qualification testing. He has a Bachelors of Science in Computer Engineering from the State University of New York at New Paltz.

A special thanks to the authors of all previous versions of this IBM Redbooks publication:

Michal Holenia
Senior Network IT Architect, IBM Slovakia

Octavian Lascu
ITSO Project Leader, IBM Romania

Rav Meerwald
Advanced Account System Service Representative, IBM Australia

Jeff Nesbitt
Advanced Account System Service Representative, IBM Australia

Pasi Laakso
IBM Certified IT Specialist, IBM Finland

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 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 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.

© Copyright International Business Machines Corporation 2015. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document REDP-5174-00 was created or updated on April 22, 2015.



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 z Systems™	System Storage®
FICON®	Parallel Sysplex®	z Systems™
GDPS®	POWER7®	z/OS®
Geographically Dispersed Parallel Sysplex™	Redbooks®	z10™
IBM®	Redpaper™	z9®
IBM z™	Redbooks (logo)  ®	zEnterprise®
	Resource Link®	

The following terms are trademarks of other companies:

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

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

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