



# IBM PureFlex Solution for Cloud Backup and Recovery: Private Cloud Disaster Recovery Strategies

## Introduction

Many organizations consider the possibility to transform their data center infrastructures in to a cloud-enabled environment by deploying private clouds. Cloud computing offers attractive opportunities to reduce costs, accelerate development, and increase the flexibility of the IT infrastructure, applications, and services. Cloud computing is the next evolutionary step that enhances traditional virtualized environments by combining server, storage, and networking resources into a single shared converged pool, and adding intellectual workload management and usage metering capabilities to the IT infrastructure.

IBM® PureFlex™ System, an expert integrated system, offers a cloud-enabled integrated server, storage, and networking platform that delivers intelligent workload deployment and scalable IT resource pools. PureFlex System increases security and resiliency to promote maximum uptime, and integrated and easy-to-use systems management reduces setup time and complexity, providing a quicker path to return on investment (ROI). Although PureFlex System provides highly resilient and flexible cloud infrastructure that supports 24x7 local operational environments, one of the key organizational business requirements is the ability to recover IT applications and data that support critical business processes quickly in case of disaster.

This IBM Redpaper™ describes disaster recovery strategies in the PureFlex System private cloud environment that is built with IBM SmartCloud® Entry running on IBM x86 and IBM POWER® compute nodes. This paper describes two-site disaster recovery approaches that use the storage replication capabilities of the IBM Flex System™ V7000 Storage Node, IBM Flex System FC5022 16Gb SAN Scalable Switches and FC adapters, IBM System Storage SAN06B-R multiprotocol routers, IBM Tivoli® Storage Manager backup and recovery technologies, and automated platform-specific disaster recovery solutions, such as VMware Site Recovery Manager or IBM PowerHA® SystemMirror® Enterprise Edition. This paper is for IT professionals who are interested in learning about common strategies for achieving disaster recovery objectives in the PureFlex System-based private cloud environment.

# IBM PureFlex System private cloud solution components

In the cloud-computing arena, the private cloud plays a fundamental role. It combines the major advantages of the public cloud, such as strong standardization, self-service automation, scalability, and metering, with the advantages of on-premise data centers.

On-premise data centers provide several advantages, such as strong security, increased customization capabilities, and increased control over quality of service (QoS).

The IBM PureFlex System private cloud solution consists of the following components:

- ▶ PureFlex System with integrated server, storage, and networking resources
- ▶ Virtualized environment with a choice of hypervisors:
  - IBM PowerVM®
  - VMware vSphere
  - Linux Kernel-based Virtual Machine (KVM)
  - Microsoft Hyper-V
- ▶ IBM SmartCloud Entry cloud management stack

## IBM PureFlex System

To meet today's complex and changing business demands, you need a solid foundation of server, storage, networking, and software resources. Furthermore, it must be simple to deploy, and able to quickly and automatically adapt to changing conditions. You must also have access to, and the ability to take advantage of, broad expertise and proven guidelines in systems management, applications, hardware maintenance, and more.

IBM PureFlex System is a comprehensive infrastructure system that provides an expert integrated computing system. It combines servers, enterprise storage, networking, virtualization, and management into a single structure. Its built-in expertise enables organizations to manage and flexibly deploy integrated patterns of virtual and hardware resources through unified management. These systems are ideally suited for customers who want a system that delivers the simplicity of an integrated solution, while still able to tune middleware and the runtime environment.

Figure 1 shows IBM PureFlex System.



Figure 1 IBM PureFlex System

IBM PureFlex System uses workload placement that is based on virtual machine (VM) compatibility and resource availability. By using built-in virtualization across servers, storage, and networking, the infrastructure system enables automated scaling of resources and true workload mobility.

IBM PureFlex System went through significant testing and experimentation so that it can mitigate IT complexity without compromising the flexibility to tune systems to the tasks that businesses demand. By offering both flexibility and simplicity, IBM PureFlex System can provide extraordinary levels of IT control, efficiency, and operating agility. With this combination, businesses can rapidly deploy IT services at a reduced cost. Moreover, the system is built on decades of expertise. This expertise enables deep integration and central management of a comprehensive and open-choice infrastructure system. It also dramatically reduces the skills and training that are required for managing and deploying the system.

The PureFlex System offers many advantages:

- ▶ Configurations that ease acquisition experience and match your needs
- ▶ Optimized to align with targeted workloads and environments
- ▶ Designed for cloud with SmartCloud Entry included on Standard and Enterprise
- ▶ Choice of architecture, operating system, and virtualization engine

- ▶ Designed for simplicity with integrated and single-system management across physical and virtual resources
- ▶ Simplified ordering that accelerates deployment in to your environments
- ▶ Sent as a single integrated entity directly to you
- ▶ Includes factory integration and lab services optimization

IBM PureFlex System combines advanced IBM hardware and software, and patterns of expertise. It integrates them into three optimized configurations that are simple to acquire and deploy so you get fast time-to-value.

The IBM PureFlex System is offered in the following configurations:

- ▶ Express, which is the infrastructure system for small-sized and mid-sized businesses, and the most cost-effective entry point.
- ▶ Standard, which is the infrastructure system for application servers with supporting storage and networking.
- ▶ Enterprise, which is the infrastructure system that is optimized for scalable cloud deployments. Enterprise has built-in redundancy for highly reliable and resilient operation to support critical applications and cloud services.

A PureFlex System configuration has these main components:

- ▶ A preinstalled and configured IBM Flex System Enterprise Chassis
- ▶ Compute nodes with either IBM POWER or Intel Xeon processors
- ▶ IBM Flex System Manager, preinstalled with management software and licenses for software activation
- ▶ IBM Storwize® V7000 external storage unit or IBM Flex System V7000 Storage Node
- ▶ All hardware components that are preinstalled in an IBM PureFlex System 42U rack
- ▶ Choice of the following items:
  - Operating system: IBM AIX®, IBM i, Microsoft Windows, Red Hat Enterprise Linux, or SUSE Linux Enterprise Server
  - Virtualization software: IBM PowerVM, Linux KVM, VMware vSphere, or Microsoft Hyper-V
  - IBM SmartCloud Entry
- ▶ Complete preintegrated software and hardware
- ▶ Onsite services that are included to get the system up and running quickly

These configurations are summarized in Table 1.

Table 1 IBM PureFlex System configurations

Component	IBM PureFlex System Express	IBM PureFlex System Standard	IBM PureFlex System Enterprise
IBM PureFlex System 42U Rack	1	1	1
IBM Flex System Enterprise Chassis	1	1	1
IBM Flex System Fabric EN4093R 10Gb Scalable Switch	1	1	2 (with both port-count upgrades)
IBM Flex System FC5022 24-port 16Gb ESB SAN Scalable Switch	1	2	2
IBM Flex System Manager Node	1	1	1
IBM Flex System Manager software license	IBM Flex System Manager (with one-year service and support)	IBM Flex System Manager Advanced (with three-year service and support)	Flex System Manager Advanced (with three-year service and support)
Chassis Management Module	2	2	2
Chassis power supplies (standard/maximum)	2/6	4/6	6/6
Chassis 80 mm fan modules (standard/maximum)	4/8	6/8	8/8
IBM Storwize V7000 Disk System or IBM Flex System V7000 Compute Node	Yes (redundant controller)	Yes (redundant controller)	Yes (redundant controller)
IBM Storwize V7000 Software	Base with one-year software maintenance agreement	Base with three-year software maintenance agreement	Base with three-year software maintenance agreement

The fundamental building blocks of IBM PureFlex System solutions are the IBM Flex System Enterprise Chassis with compute nodes, networking, and storage.

## IBM Flex System

IBM Flex System is built from reliable IBM technology that supports open standards and is designed for multiple generations of technology. This enables clients to support existing workloads while being ready for the future demands of their business. Starting with an innovative chassis design for new levels of simplicity, reliability, and upgradability, IBM Flex System helps clients go beyond blade servers. With a broad range of x86 and IBM POWER compute nodes, the Flex System V7000 storage node, enhanced networking capabilities, and sophisticated system management capabilities, clients can upgrade their existing blade server infrastructure to make IT simpler, more flexible, more open, and more efficient.

IBM Flex System is shown in Figure 2.



Figure 2 IBM Flex System

## Management

IBM Flex System Manager is designed to optimize the physical and virtual resources of the IBM Flex System infrastructure while simplifying and automating repetitive tasks. It provides easy system setup procedures with wizards and built-in expertise, and consolidated monitoring for all of your resources, including compute, storage, networking, virtualization, and energy. IBM Flex System Manager provides core management functionality along with automation. It is an ideal solution that allows you to reduce administrative expense and focus your efforts on business innovation.

A single user interface controls the following features:

- ▶ Intelligent automation
- ▶ Resource pooling
- ▶ Improved resource utilization
- ▶ Complete management integration
- ▶ Simplified setup

## Compute nodes

The compute nodes take advantage of the full capabilities of IBM POWER7® and Intel Xeon processors. This configuration offers the performance that you need for your critical applications.

With support for a range of hypervisors, operating systems, and virtualization environments, the compute nodes provide the foundation for the following environments:

- ▶ Virtualization solutions
- ▶ Database applications
- ▶ Infrastructure support
- ▶ Line-of-business applications

## Storage

The storage capabilities of IBM Flex System V7000 storage node give you advanced functionality with storage nodes in your system, and take advantage of your existing storage infrastructure through advanced virtualization.

IBM Flex System V7000 storage node simplifies storage administration with a single user interface for all your storage. The management console is integrated with the comprehensive management system. These management and storage capabilities allow you to virtualize third-party storage with nondisruptive migration of your current storage infrastructure. You can also take advantage of intelligent tiering so you can balance performance and cost for your storage needs. The solution also supports local and remote replication, and snapshots for flexible business continuity and disaster recovery capabilities.

## Networking

The range of available adapters and switches to support key network protocols allow you to configure IBM Flex System to fit in your infrastructure. However, you can do so without sacrificing being ready for the future. The networking resources in IBM Flex System are standards-based, flexible, and fully integrated into the system. This combination gives you no-compromise networking for your solution. Network resources are virtualized and managed by workload. These capabilities are automated and optimized to make your network more reliable and simpler to manage.

IBM Flex System gives you the following key networking capabilities:

- ▶ Supports the networking infrastructure that you have today, including Ethernet, Fibre Channel, and InfiniBand.
- ▶ Offers industry-leading performance with 1 Gb and 10 Gb Ethernet, 8 Gb and 16 Gb Fibre Channel, and FDR InfiniBand.
- ▶ Provides pay-as-you-grow scalability so you can add ports and bandwidth when needed.

## Enterprise Chassis

The IBM Flex System Enterprise Chassis is the foundation of the offering, supporting intelligent workload deployment and management for maximum business agility. The 14-node, 10U chassis delivers high-performance connectivity for your integrated compute, storage, networking, and management resources. The chassis is designed to support multiple generations of technology, and offers independently scalable resource pools for higher utilization and lower cost per workload.

## PureFlex System and Flex System as a cloud platform

Fast deployment, simple to manage, easily scalable as workload grows, lower total cost of ownership, security, and resiliency are the key attributes that make the IBM PureFlex System a great choice for cloud deployments. IBM Flex System, as a building block of IBM PureFlex System, offers highly customizable and flexible component selection to meet specific company needs.

Table 2 summarizes the attributes of IBM PureFlex System and IBM Flex System.

Table 2 IBM PureFlex System and IBM Flex Systems attributes

Attribute	IBM PureFlex System	IBM Flex System
Flexibility	Predefined subset of the solution components.	The entire set of components is available for selection.
Customization	Only selected components are configurable.	The entire set of solution components can be configured.
Deployment	Rack-level preintegrated and factory build; IBM installation services are included to help with the onsite deployment.	Chassis-level and compute node-level factory integration; onsite assembly by the client or IBM Business Partner.
Management	Has centralized management with unified tools.	
Scalability	Scales elastically to meet workload demands.	
Availability and resiliency	Is designed for continuous business operations.	
Security	Maintains integrity and privacy.	
Energy efficiency	Has shared architecture with energy efficient components.	

IBM PureFlex System and IBM Flex System based cloud solutions can provide a cost-efficient, secure, and scalable platform for diverse workloads, which can offer the following benefits:

- ▶ Accelerate time-to-value: Accelerate deployment, speed response to changing business needs, and improve productivity with a range of pre-integrated/configured and optimized infrastructure solutions.
- ▶ Maximize IT availability: Improve overall customer service and satisfaction with a cloud infrastructure that is designed for continuous operation.
- ▶ Scale elastically with better provisioning of resources: Enhance service quality and delivery with innovative technologies and scalable offerings, which enable highly dynamic and fluid, just-in-time and online resource configuration to optimize cloud delivery for a wider-range of workloads.
- ▶ Ensure security: Deliver a comprehensive security solution to mitigate risk by maintaining integrity and privacy of sensitive services and data in a cloud across the entire stack.
- ▶ Reduce IT costs: Reduce overall costs and improve the ROI of IT for cloud delivery with integrated solutions that enable consolidation, better usage, and better and easier management of capital assets.

## IBM SmartCloud Entry

IBM SmartCloud is a family of integrated enterprise-class cloud-computing technologies and services for securely building and using private, public, and hybrid clouds. Built upon open standards that are combined with IBM expertise and best practices, SmartCloud is designed for complex high-performance computing environments. SmartCloud goes beyond securely delivering new cloud efficiencies and cost savings to drive fundamental innovation for lasting marketplace advantage.

For more information about IBM SmartCloud, go to the following website:

<http://www.ibm.com/cloud-computing/>



IBM SmartCloud Entry is robust software that simplifies the management of your virtualized environment. Users can request and provision an environment quickly through an easy-to-use web-based interface. IT managers can monitor and manage this environment for improved efficiency and utilization of the data center. IBM SmartCloud Entry includes heterogeneous server architecture support from a single SmartCloud Entry user interface. This support enables greater flexibility, as customers can manage their environment across IBM System x®, BladeCenter®, Flex System, and PureFlex System platforms with a single, consolidated image of SmartCloud Entry software.

By using IBM SmartCloud Entry, you can maintain control over the allocation of resources with a web-based application. You can perform the following common public or private cloud operations:

- ▶ Provisioning and de-provisioning virtual servers
- ▶ Drafting and cloning workloads
- ▶ Capturing workloads
- ▶ Starting and stopping virtual servers as part of a workload
- ▶ Resizing existing virtual servers
- ▶ Creating projects to give team-specific access to workloads
- ▶ Providing network configurations that set unique network properties to different workloads
- ▶ Billing, accounting, and metering support
- ▶ Providing request and approval workflow support

An IBM SmartCloud Entry environment consists of a virtual machine installation that integrates with the virtualization management software and manage virtualized IBM x86 and IBM Power compute nodes:

- ▶ IBM x86 compute nodes: VMware vSphere hypervisors (ESXi) that are managed by vCenter.
- ▶ IBM Power compute nodes: PowerVM hypervisor, VIOS, and IBM Flex System Manager™ with activated VMControl Enterprise Edition.

## IBM PureFlex System disaster recovery concepts

Disaster recovery (DR) is the ability to restore IT operations of a data center at a *different* location on *different* hardware if the primary location experiences a catastrophic failure and becomes inoperable. Disaster recovery procedures are documented in the disaster recovery plan, and they are periodically tested to ensure that everything is working as designed. A disaster recovery plan is part of an organization's overall business continuity strategy.

The goal of any disaster recovery solution is to protect the most business critical processes and data from outages, and minimize unplanned downtime. All planning for any type of DR solution is always subject to balancing the solution's downtime goals with its costs. The following sections define terms and a structure for DR solutions to describe and quantify their characteristics.

In a disaster situation, users are normally aware that an outage occurred. Two of the dominating properties of a proposed DR solution are based on the user visible duration of the computing services impact. The properties are as follows:

- ▶ Recovery Point Objective (RPO)
- ▶ Recovery Time Objective (RTO)

The *Recovery Point Objective* (RPO) represents the extent of data loss you are willing to accept because of a disaster. It is measured as the duration of time before the disaster event for which you must rerun your work (or accept its loss) after your system is recovered.

The RPO is a requirement for the currency of data that is available for recovery. For example, a business that believes it could acceptably afford to re-create or lose the data that was processed in the last five minutes preceding a disaster event has an RPO of five minutes. A business that relies on daily backups has to plan for a 24 hour RPO for a site disaster.

The *Recovery Time Objective* (RTO) is the duration of time that follows a disaster for which you are willing to accept the loss of computing services. The period starts from the moment of the disaster until the moment when the systems are recovered. You can consider the RTO as a measure of how long a business can afford to have systems and applications down after a disaster. For example, a business that believes that it could afford to be without systems for eight hours has an RTO of eight hours.

For this paper, we focus on RPO and RTO that can be achieved by using the PureFlex System integrated replication capabilities that are available in IBM Flex System V7000 Storage Node, which are combined with a disaster recovery-capable software, such as Tivoli Storage Manager or VMware Site Recovery Manager.

## IBM Flex System V7000 replication family services

The term *Replication Family Services* refers to the various copy services that are available on the IBM Flex System V7000. These copy services include the following ones:

- ▶ IBM FlashCopy®
- ▶ IBM Metro Mirror and Global Mirror
- ▶ IBM Global Mirror with Change Volumes

### FlashCopy

FlashCopy is a function that allows you to create a point-in-time copy of one of your Flex System V7000 disks. This function might be helpful when you perform backups or application testing. These copies may be cascaded upon one another, read from, written to, and even reversed.

These copies can conserve storage, if needed, by being space-efficient copies that only record items that changed from the originals instead of full copies.

### Metro Mirror and Global Mirror

Metro Mirror and Global Mirror are technologies that enable you to keep a real-time copy of a disk at a remote site that contains another Flex System V7000 system.

*Metro Mirror* is designed for metropolitan distances with a zero RPO, that is, zero data loss. This is achieved with a synchronous copy of volumes; writes are not acknowledged until they are committed to both storage systems. By definition, synchronous replication makes the host wait for write I/Os to complete at both the local and remote storage systems and includes round-trip network latencies. Metro Mirror has the following characteristics:

- ▶ Zero RPO
- ▶ Synchronous
- ▶ Production application performance that is impacted by round-trip latency

*Global Mirror* is designed to minimize application performance impact by replicating asynchronously, that is, writes are acknowledged as soon as they can be committed to the local storage system, sequence-tagged, and passed on to the replication network. This allows Global Mirror to be used over longer distances. By definition, asynchronous replication results in an RPO greater than zero. However, in the case of Global Mirror, the RPO is small, typically anywhere from several milliseconds to a few seconds.

Although Global Mirror is asynchronous, the RPO is still small and thus the network and the remote storage system must both still be able to cope with peaks in traffic. Global Mirror has the following characteristics:

- ▶ Near-zero RPO
- ▶ Asynchronous
- ▶ Production application performance that is impacted by I/O sequencing preparation time

### **Global Mirror with Change Volumes**

This function (also known as Cycle-Mode Global Mirror) can best be described as “Continuous Remote FlashCopy.” If you use this feature, the Flex System V7000 essentially takes periodic FlashCopy snapshots of a disk and writes them to your remote destination.

*Global Mirror with Change Volumes* provides an option to replicate point-in-time copies of volumes. This generally requires lower bandwidth because it is the average rather than the peak throughput that must be accommodated. The RPO for Global Mirror with Change Volumes is higher than traditional Global Mirror. Global Mirror with Change Volumes has the following characteristics:

- ▶ Larger RPO
- ▶ Point-in-time copies
- ▶ Asynchronous
- ▶ Possible system performance impact because point-in-time copies are created locally

## **IBM PureFlex System disaster recovery topologies**

Remote replication services require an underlying network infrastructure that can provide the required bandwidth and latency to meet target RPO. In general, there are two types of connectivity that can be established between primary and recovery sites:

- ▶ Dark fiber/Wave Division Multiplexing (WDM)
- ▶ Fibre Channel over IP (FCIP) tunneling

### **Dark fiber/WDM**

Dark fiber/WDM provides high-bandwidth low-latency connectivity that is suitable for synchronous replication to ensure that the application response times are not affected. However, the cost that is associated with the maintenance of these lines is typically high, and the distances between sites are limited to a few hundred kilometers.

With WDM, both local and remote sites have a common Fibre Channel fabric, and the storage devices that are attached to this fabric are accessible from both sites.

Figure 3 shows a typical WDM-based inter-site connectivity architecture.

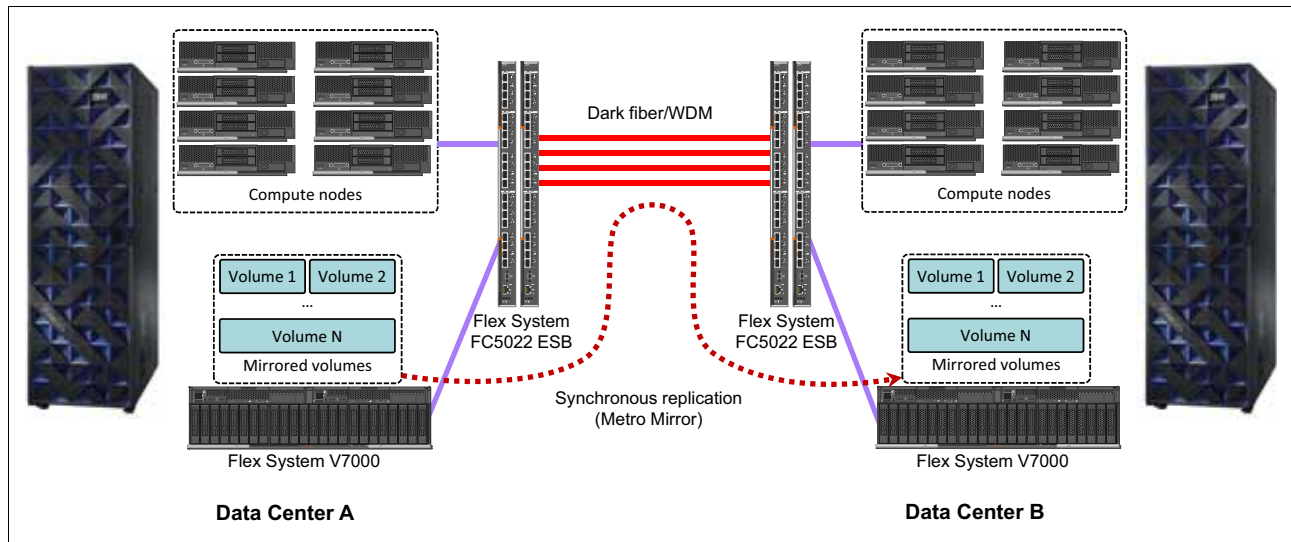


Figure 3 Dark fiber/WDM inter-site connectivity architecture

The key component of the PureFlex System data center that is used for WDM-based connectivity is IBM Flex System FC5022 24-port 16Gb ESB Scalable Switch.

The 24-port Enterprise Services Bundle (ESB) model is a fully enabled switch with a complete set of licensed features that maximizes performance, ensures availability, and simplifies management for the most demanding applications and expanding virtualization environments. The factory-installed licenses include the following items:

- ▶ ISL Trunking: Consolidates ISLs into fault-tolerant and load-balanced trunks with bandwidth of up to 128 Gbps for greater performance and simplified management.
- ▶ Adaptive Networking: Ensures high-priority connections that obtain the bandwidth that is necessary for optimum performance.
- ▶ Advanced Performance Monitoring: Provides end-to-end visibility into fabric performance to maximize performance tuning and the utilization of fabric resources.
- ▶ Fabric Watch: Monitors and creates alerts that are based on the health of switch and fabric elements.
- ▶ Extended Fabrics: Extends Fibre Channel SANs beyond traditional distance limitations for replication and backup at full bandwidth.
- ▶ Server Application Optimization: Allows virtual machine quality of service, authentication, and monitoring from SAN through an adapter to virtual machine and associated virtualized application.

An Extended Fabrics license allows these switches to interact over the long distances (typically up to 300 km) to support metropolitan area disaster recovery configurations.

The ISL Trunking feature can also be used in a WDM scenario to increase the available bandwidth between sites to satisfy the requirements of the most demanding workloads being replicated.

FC5022 Scalable Switch also supports ClearLink diagnostic tests on ISL links. ClearLink ensures signal integrity between the sites, and it also allows you to test the latency of the link and the full line rate capability based on a predefined traffic pattern.

Inter-site communication security and bandwidth usage questions can be addressed by using in-flight encryption and compression at line rate speeds that are supported by F5022 switches.

## FCIP

FCIP provides alternative, lower-cost method of connecting two sites for disaster recovery purposes, at the expense of lower bandwidth and higher latency. FCIP also supports extended distances, potentially over few thousands kilometers, and it is most commonly used with asynchronous replication.

Figure 4 shows a typical FCIP-based inter-site connectivity.

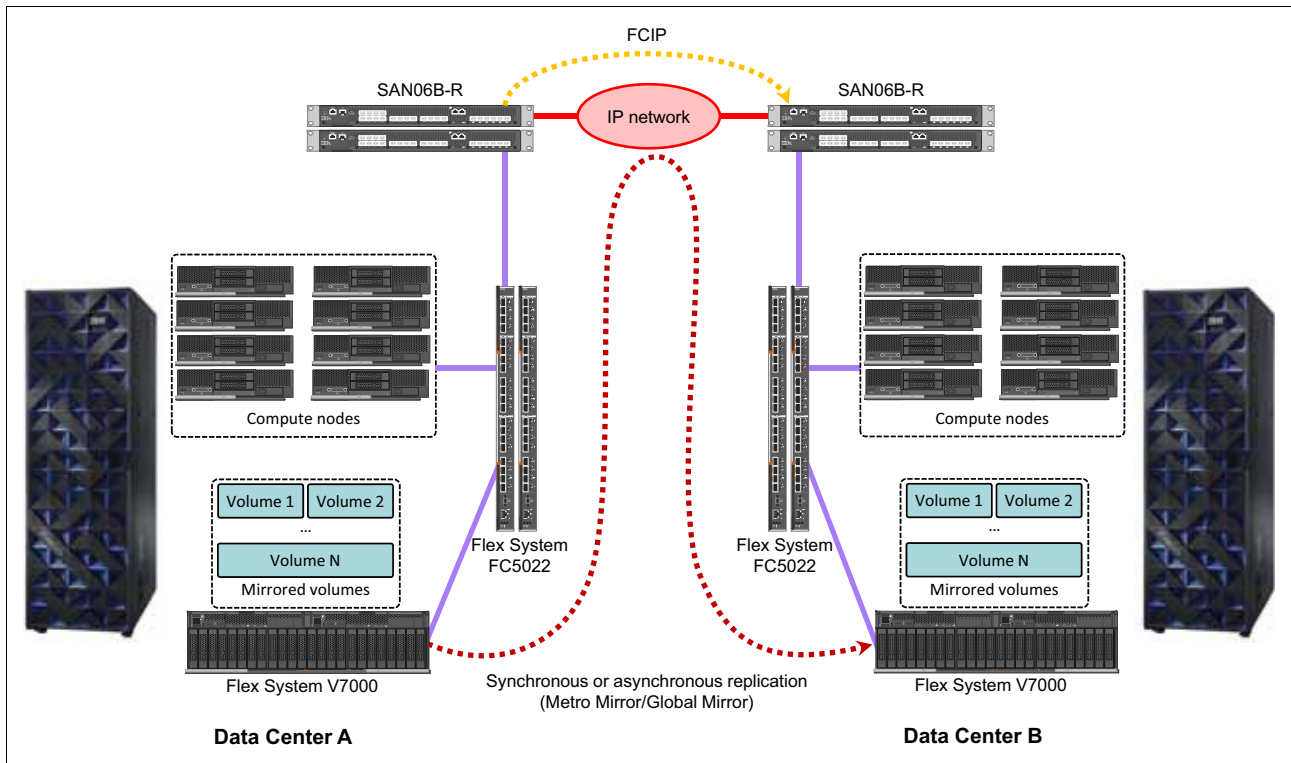


Figure 4 FCIP inter-site connectivity architecture

A local PureFlex System SAN fabric is built with the FC5022 switches (ESB models are not required), and the multiprotocol routers, such as IBM System Storage® SAN06B-R.

The multiprotocol router provides Fibre Channel FC-FC Routing Service, which allows the interconnection of multiple SAN islands without requiring that the separate fabrics be merged into a single large SAN. It can help create a tiered or extended enterprise SAN infrastructure without redesigning or reconfiguring the entire environment.

The IBM System Storage SAN06B-R multiprotocol router provides Fibre Channel over IP and FCIP Tunneling Service for distance extension, which can enable cost-effective and manageable metro and global business continuity solutions. This extended distance connectivity can help create consolidated remote tape vaulting data protection plus Metro Mirror and Global Mirror disk-based, disaster-tolerant solutions.

The IBM SAN06B-R multiprotocol router has the following features:

- ▶ Designed for high performance with 8 Gigabit per second (Gbps) Fibre Channel (FC) ports and hardware-assisted traffic processing for line-rate performance across Ethernet IP ports.
- ▶ Uses existing Internet IP-based Metropolitan Area Network (MAN) or wide area network (WAN) infrastructures for metro and global SAN extension for business continuity solutions.
- ▶ Enables consolidation of storage area network (SAN) islands for infrastructure simplification without compromising security.
- ▶ Hardware-based compression, large window sizes, and selective acknowledgement of IP packets optimize the performance of SAN extension over IP networks.
- ▶ IPSec security protects all TCP traffic over an FCIP tunnel to establish secure communications between the sites.
- ▶ Eight virtual FCIP tunnels per IP port are enabled by the High-Performance Extension feature to help maximize scalability and utilization of MAN/WAN resources.
- ▶ Integrated IBM System Storage SAN b-type switch management helps simplify installation and administration and helps provide fabric investment protection.

IBM Flex System FC5022 switches (non-ESB) come standard with the following software features:

- ▶ Full Fabric mode: Enables high performance 16 Gb or 8 Gb fabric switching.
- ▶ Access Gateway mode: Uses NPIV to connect to any fabric without adding switch domains to reduce management complexity.
- ▶ Dynamic Path Selection: Enables exchange-based load balancing across multiple Inter-Switch Links for superior performance
- ▶ Advanced Zoning: Segments a SAN into virtual private SANs to increase security and availability.
- ▶ Enhanced Group Management: Enables centralized and simplified management of SAN fabrics through IBM Network Advisor.

Optional ISL Trunking and Fabric Watch features can be activated on FC5022 switches to increase inter-switch link (ISL) bandwidth and to provide advanced monitoring capabilities in the fabric.

Table 3 summarizes the characteristics of WDM and FCIP-based inter-site connectivity.

Table 3 WDM and FCIP comparison

Feature	WDM	FCIP
Distance	Up to 300 km	Up to 1000's km
Replication type	Synchronous	Asynchronous or synchronous
Networking hardware	SAN switch with Extended Fabrics license	SAN switch and SAN multiprotocol router
Disaster recovery type	Highly automated business continuity solutions (RTO units of measure: minutes)	<ul style="list-style-type: none"> <li>▶ Asynchronous replication: Backup/rapid restore (RTO units of measure: hours)</li> <li>▶ Synchronous replication: Highly automated business continuity solutions (RTO units of measure: minutes)</li> </ul>
Relative technology cost	High	Low

# IBM PureFlex System disaster recovery solutions

In general, there are three types of business continuity and disaster recovery approaches:

- ▶ Core storage technologies

Core technologies include advanced copy services that enable reliability, redundancy, and business continuity, at the solution and system operations level (rather than at the component level).

Core technologies can be in server operating systems, storage systems, storage software, file systems, and other components. They also include management software and services for enterprise-wide policy-based management of backup copies and administration of the storage replication services.

- ▶ Platform-specific integration

Each operating system platform and virtualized environment has its own requirements for how to access and use the data, volume copies, and replication recovery copies. This layer provides operating system- or environment-specific tasks and integration that are needed to use the copied data.

- ▶ Application-specific integration

Applications need to have specific integration into the previous layers to take full advantage of the underlying functions. This takes the form of integrated application commands, familiar to the application users, that can transparently use advanced server or storage capabilities. This functionality often includes coordinating core technology operations across many multiple LUNs, data sets, and objects, disk systems, and servers that might make up an application system.

This paper describes three disaster recovery solutions that use two disaster recovery approaches:

- ▶ Disaster recovery with Tivoli Storage Manager (core storage technologies)
- ▶ Disaster recovery with VMware Site Recovery Manager (platform-specific integration)
- ▶ Disaster recovery with IBM PowerHA SystemMirror Enterprise Edition (platform-specific integration)

## Disaster recovery with Tivoli Storage Manager

One of the common methods for achieving disaster recovery objectives is an automated backup solution with rapid recovery. A PureFlex System private cloud is a virtualized environment, and backups in virtualized environments can be separated into in-guest backup and off-host backup types.

The backup types that are used in virtual environments include the following ones:

- ▶ In-guest backups are started in the virtual machine and provide the following features:
  - The backup agent is installed on the guest operating system.
  - A file-level backup or a block-level image backup can be created and recovered.
  - The backup model is similar to backups on physical hosts.
  - For Tivoli Storage Manager backups, the Backup-Archive client, the data protection clients, or IBM Tivoli Storage Manager FastBack® is used.
  - Protection is provided for databases and applications.

- ▶ Off-host backups are started on a dedicated backup server and provide the following features:
  - The backup agent can be installed on the backup server (such as vStorage Backup Server in VMware environments).
  - File-level and guest-level image backups can be created and recovered.
  - Centralized management of backup data is provided.
  - Backups are offloaded from the host to another host to free the production server resources.
  - File system consistent backups can be created by using snapshots.

IBM offers IBM Tivoli Storage Manager, a suite of backup and recovery software tools that can be used for backup and disaster recovery purposes in a PureFlex System cloud environment.

Tivoli Storage Manager is a storage management application that is built for the enterprise. Tivoli Storage Manager provides an enterprise solution for data protection, disaster recovery, space management, and record retention. Tivoli Storage Manager facilitates flexible and scalable storage management policies to support complicated business needs for storage management and disaster recovery. Most importantly, Tivoli Storage Manager automates storage management tasks by eliminating labor and cost intensive manual procedures for backup, archive, and recovery.

Tivoli Storage Manager includes many components that are targeted to perform specific backup and recovery functions for different applications and data, and for virtualized environments. For complex environments, IBM offers IBM Tivoli Storage Manager Suite for Unified Recovery, a pre-packaged suite of Tivoli Storage Manager software products with simplified ordering and licensing.

### **Tivoli Storage Manager Suite for Unified Recovery**

Tivoli Storage Manager Suite for Unified Recovery is a bundle of ten proven data protection and recovery software products (see Figure 5 on page 18). This bundle helps organizations meet a wide range of data management challenges for complex, distributed infrastructures. You can deploy the advanced management tools that you need for each of your individual data protection requirements without having to worry about individual product licenses.

Tivoli Storage Manager Suite for Unified Recovery offers the following features:

- ▶ Provides extensive data protection for a wide range of systems, including VMs, file servers, email, databases, mainframes, and even desktops. This bundled solution allows you to use the correct data protection tool for each of your requirements.
- ▶ Reduces costs and simplifies procurement and deployment with per-terabyte capacity licensing. You can deploy any of 10 separate solution components, in any location and quantity, with a simplified license that measures only the amount of data that is managed.
- ▶ Scales to meet the recovery needs of any size organization by managing up to four billion data objects on a single server. This solution supports more than 50 operating system versions and hundreds of server and storage devices.
- ▶ Manages the entire suite of products from a single user console. You can configure, manage, upgrade, report, and monitor all ten products from a single administration interface.



IBM Tivoli Storage Manager Suite for Unified Recovery is packaged to address the following challenges:

- ▶ Various technologies are needed to protect various types of systems, applications, and data in various locations, to help meet the service level requirements of the business.
- ▶ Procurement and management of multiple data recovery point solutions from separate vendors, or even multiple add-on products from the same vendor, can cause stress among your IT staff and introduce additional risk of human error.

Tivoli Storage Manager Suite for Unified Recovery provides a full suite of products with a licensing model that is easy to measure:

- ▶ Centralized and simplified management: One user interface (UI) provides access to control and monitor a wide range of data backup and recovery processes across the enterprise.
- ▶ Simplified procurement: One part number provides access to unlimited usage of 10 separate solution components
- ▶ Easier and more straightforward budget control: Predicting future requirements can be easier.

Tivoli Storage Manager Suite for Unified Recovery includes 10 proven data protection and recovery software products. Each component can be deployed as needed in any quantity to meet specific service-level requirements with no per-product licensing charges.

This advanced and highly scalable suite helps increase the efficiency of your IT operations and helps cut costs that are related to storage management. These improved efficiencies and cost-cuts are done by providing a wide range of data protection, recovery management, and monitoring capabilities by using policy-based automation, including the following items:

- ▶ Backup and recovery, and archive and retrieval
- ▶ Snapshots for online database and application protection
- ▶ Disaster recovery planning and replication
- ▶ Bare machine recovery
- ▶ Data deduplication and space management

The components of the suite are shown in Figure 5.

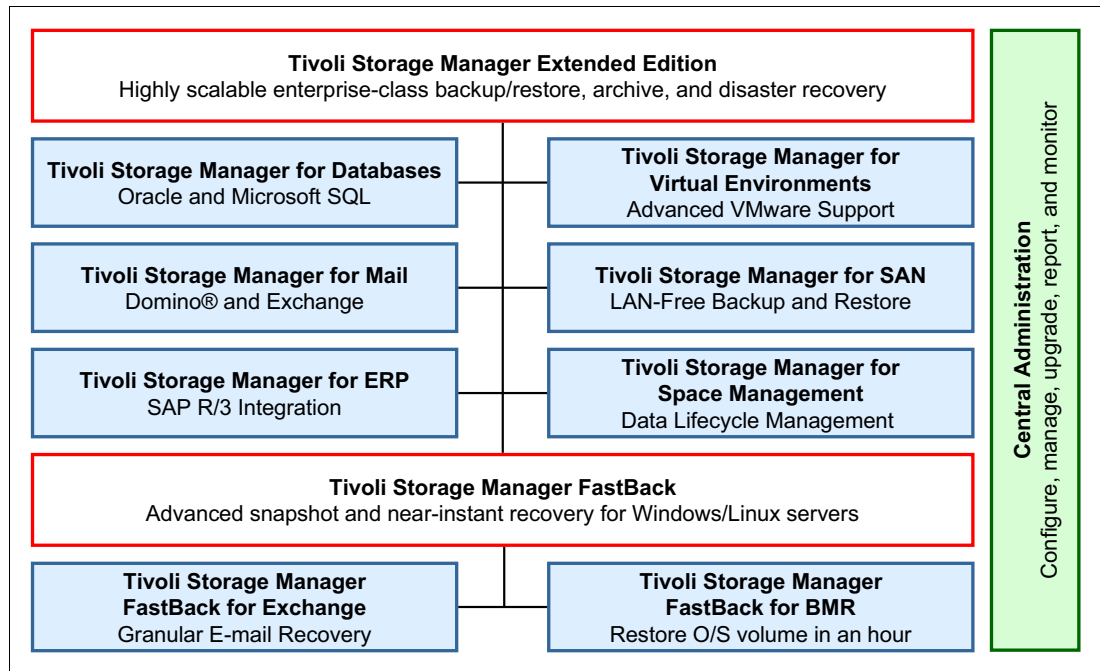


Figure 5 Tivoli Storage Manager Suite for Unified Recovery

The individual products that are included in this comprehensive package are as follows:

- ▶ **Tivoli Storage Manager Extended Edition**  
Provides core backup and restore operations for a wide range of operating systems, broad support for tiers of storage, network data management protocol (NDMP), IBM DB2® and IBM Informix® support, source and target deduplication, and advanced disaster recovery planning and replication.
- ▶ **Tivoli Storage Manager for Databases**  
Performs online, consistent, and centralized backups for Oracle and SQL to avoid downtime, protect vital enterprise data infrastructures, and minimize operation costs.
- ▶ **Tivoli Storage Manager for Enterprise Resource Planning**  
Performs online, consistent, and centralized backups for SAP environments.
- ▶ **Tivoli Storage Manager for Mail**  
Protects data on email servers that are running IBM Lotus® Domino® or Microsoft Exchange, with granular restoring of Exchange email objects.
- ▶ **Tivoli Storage Manager for Virtual Environments**  
Automatically discovers and protects VMware virtual machines, offloads backup workloads to a centralized server, and enables flexible, near-instant recovery.
- ▶ **Tivoli Storage Manager for Space Management**  
Provides hierarchical storage management (HSM) to automatically migrate rarely accessed files to alternative storage, without disrupting the most frequently used files in local storage. Migrated files are automatically and transparently recalled to primary storage when required by applications or users.

- ▶ Tivoli Storage Manager for Storage Area Networks
 

Provides high-performance backup and restore operations by removing data transfer from the local area network (LAN).
- ▶ IBM Tivoli Storage Manager FastBack
 

Provides efficient block-level incremental backup and near-instant restore operations for critical Microsoft Windows and Linux servers and applications, both in the data center and in remote offices.
- ▶ Tivoli Storage Manager FastBack for Microsoft Exchange
 

Enables the recovery of individual Microsoft Exchange objects such as email, calendar entries, attachments, contacts, and tasks.
- ▶ Tivoli Storage Manager FastBack for Bare Machine Recovery
 

Provides operating system volume recovery following a disaster or catastrophic server failure, fully restoring Windows and Linux systems within an hour.

Tivoli Storage Manager Suite for Unified Recovery delivers centralized administration and intelligent data move-and-store techniques to help ease storage management. The product scales from small to large installations, managing up to four billion data objects in a single server. It supports more than 50 operating system versions and hundreds of server and storage devices. It also facilitates a multitude of connections, including Internet, wide area networks (WANs), LANs, and storage area networks (SANs).

Optionally, IBM Tivoli Storage FlashCopy Manager can be used for performing and managing frequent, near-instant, nondisruptive, and application-aware backups and restores that use advanced FlashCopy snapshot technologies in IBM storage systems, including Flex System V7000.

## **Tivoli Storage FlashCopy Manager**

As the amount of data that needs protecting continues to grow exponentially and the need to keep the downtime that is associated with backup to an absolute minimum, Tivoli Storage FlashCopy Manager can help minimize the impact that is caused by backups and provide near instant restore capabilities with support for data volume snapshot technologies, such as FlashCopy or Microsoft Volume Shadow Copy Services (VSS).

Although many storage systems are now equipped with volume snapshot tools, these hardware-based snapshot technologies provide only “crash consistent” copies of data. Many business critical applications, including those that rely on a relational database, need an additional snapshot process to ensure that all parts of a data transaction are flushed from memory and committed to disk before the snapshot to ensure that you have a usable and consistent copy of the data.

Tivoli Storage FlashCopy Manager helps deliver the highest levels of protection for mission-critical IBM DB2, SAP, Oracle, Microsoft Exchange, and Microsoft SQL Server applications through integrated and application-aware snapshot backup and restore capabilities. This is achieved through the exploitation of advanced IBM storage hardware snapshot technology to create a high performance and low impact application data protection solution.

The snapshots that are captured by Tivoli Storage FlashCopy Manager can be retained as backups on local disk, and with optional integration with Tivoli Storage Manager, customers can use the full range of advanced data protection and data reduction capabilities, such as data deduplication, progressive incremental backup, hierarchical storage management, and centrally managed policy-based administration.

Because a snapshot operation typically takes much less time than the time for a tape backup, the window during which the application must be aware of a backup can be reduced. This facilitates more frequent backups, which can reduce the time that is spent performing forward recovery through transaction logs, increases the flexibility of backup scheduling, and eases administration.

Application availability is also improved because of the reduction of the load on the production servers. Tivoli Storage FlashCopy Manager uses storage snapshot capabilities to provide high speed, low impact, and application-integrated backup and restore functionality for the supported application and storage environments. Automated policy-based management of multiple snapshot backup versions, together with a simple and guided installation and configuration process, provide an easy to use and quick to deploy data protection solution that enables the most stringent database recovery time requirements to be met.

For more information about IBM Tivoli Storage FlashCopy Manager, go to the following website:

<https://www-01.ibm.com/software/tivoli/products/storage-flashcopy-mgr/>

### Disaster recovery topology

Figure 6 shows a typical PureFlex System cloud disaster recovery architecture with Tivoli Storage Manager and replication services.

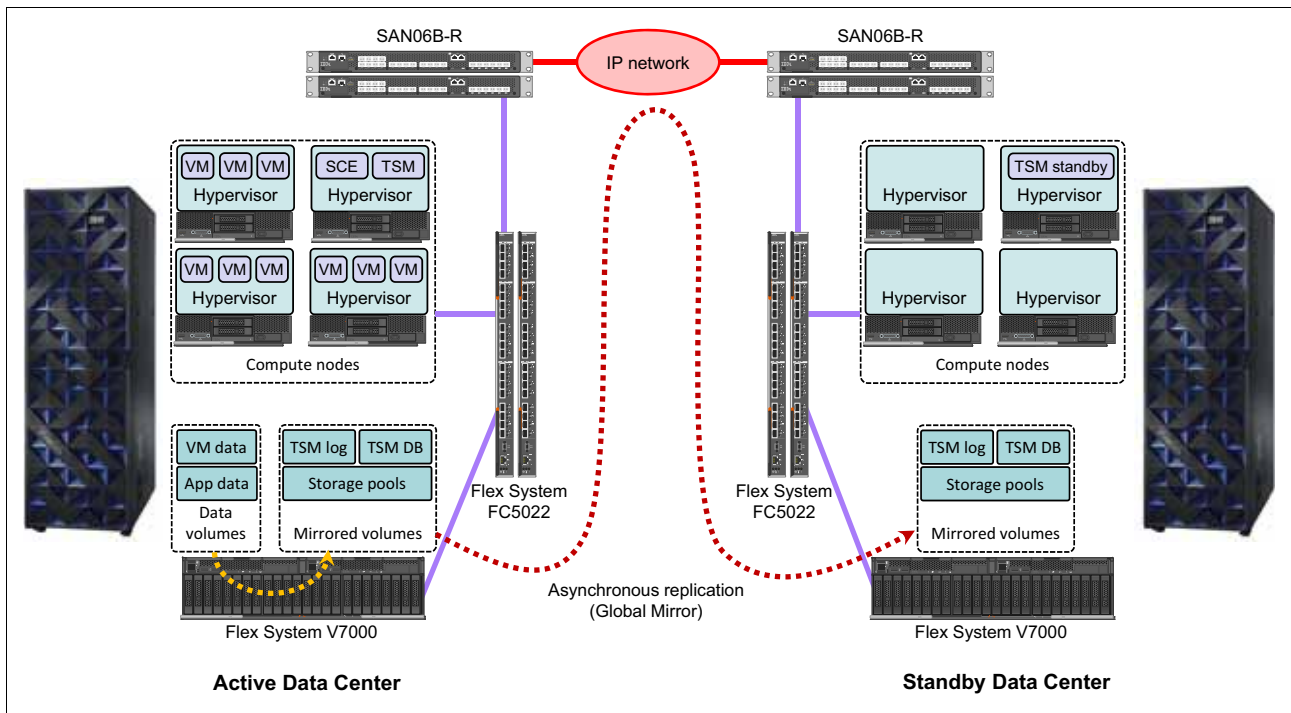


Figure 6 Disaster recovery architecture with Tivoli Storage Manager and replication services

Two PureFlex System-based data centers are located at two different sites. The data centers operate in active/standby mode, where the active data center supports the production environment, and the standby data center maintains the data that is required to successfully restore IT operations in case of a catastrophic failure at the active data center. Storage replication connectivity is based on FCIP topology (see “FCIP” on page 13).

The PureFlex System Tivoli Storage Manager disaster recovery architecture includes the following functional elements:

- ▶ **Active Tivoli Storage Manager server**

An Active Tivoli Storage Manager server operates in the active data center, and it performs scheduled policy-based backup operations of virtual machines, applications, and data.
- ▶ **Tivoli Storage Manager databases and logs**

Tivoli Storage Manager uses databases and logs to store metadata, to track its backup/recovery operations, and to recover Tivoli Storage Manager operations in case of a Tivoli Storage Manager server failure. Tivoli Storage Manager database and logs are asynchronously replicated to the standby site through Global Mirror.
- ▶ **Tivoli Storage Manager storage pools**

Tivoli Storage Manager disk-based storage pools are used to store backups of VMs, applications, and data. Storage pools are replicated to the standby site, and they are used to restore VMs, applications, and data from backup. SmartCloud Entry virtual machine is also backed up by Tivoli Storage Manager.
- ▶ **Standby Tivoli Storage Manager server**

A Standby Tivoli Storage Manager server is located at the standby data center, and it does not perform any operations. It is used only in case of a disaster at the primary location to restore normal Tivoli Storage Manager operations at the standby site by using replicated Tivoli Storage Manager databases and log files.

Backup operations are almost always automated by using Tivoli Storage Manager scheduling facilities, but recovery operations almost always occur *ad hoc* and are manually initiated by either the user or administrator initiated commands. For DR of enterprise data, recovery priorities and the sequence of data recovery events must be defined in the DR plan.

In the described example, the following steps can be performed to restore IT operations at the standby site:

- ▶ Make all mirrored volumes usable by breaking mirroring pairs.
- ▶ Restore Tivoli Storage Manager server at the standby site by using replicated Tivoli Storage Manager database and log volumes.
- ▶ Restore critical workloads by using Tivoli Storage Manager recovery from mirrored storage pools.
- ▶ Restore other workloads by using Tivoli Storage Manager recovery from mirrored storage pools.
- ▶ Restore SmartCloud Entry management functions.

## Disaster recovery with VMware Site Recovery Manager

VMware offers *Site Recovery Manager* (SRM) as an optional element of VMware Infrastructure (VI) to enable automated disaster recovery for virtual environments. Protection can extend from individual replicated data stores to an entire virtual site. The VMware virtualization of the data center offers advantages that can be applied to business continuity and disaster recovery.

The entire state of a virtual machine (memory, disk images, I/O, and device state) is encapsulated. Encapsulation enables the state of a virtual machine to be saved to a file. Saving the state of a virtual machine to a file allows the transfer of an entire virtual machine to another host.

vCenter Site Recovery Manager uses array-based replication between a primary site and recovery site, such as the Metro Mirror functionality. The workflow that is built into SRM automatically discovers which data stores are set up for replication between the protected and recovery sites. SRM can be configured to support bidirectional protection between two sites.

vCenter Site Recovery Manager provides protection for the operating systems and applications that are encapsulated by the virtual machines that are running on VMware vSphere ESX Servers.

SRM uses the IBM SAN Volume Controller *Storage Resource Adapter* (SRA) to enable the management of advanced copy services on Flex System V7000, such as Metro Mirror and Global Mirror, and is based on the Storage Networking Industry Association (SNIA) SMI-S interface.

The benefit of adding the SRA to SRM is that it enables VMware to manage the storage's remote copy functions as well as the ESX hosts and guest VMs. It aids in the following tasks:

- ▶ Setting up the recovery infrastructure
- ▶ Creation of recovery plans
- ▶ Testing recovery plans
- ▶ Automating failover

Figure 7 shows a high-level overview of the relationships between vCenter, SRM, the ESX hosts and VMs, and the storage at the local and remote sites.

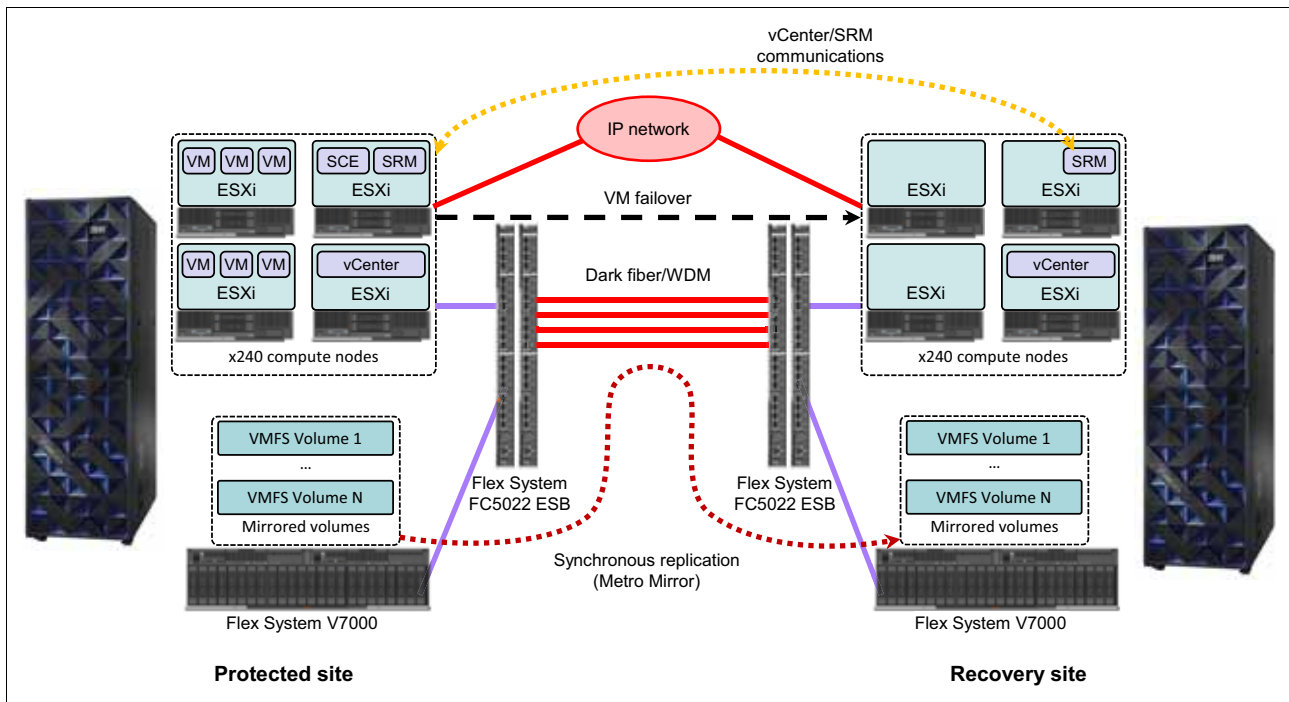


Figure 7 VMware Site Recovery Manager topology

In the example that is shown in Figure 7, vCenter and SRM servers are set up on both protected and recovery sites, and the Metro Mirror synchronous replication over the WDM is used for mirroring VMFS partitions.

In case of a disaster at the protected site, SRM assists in performing VM failover to the recovery site by restarting virtual machines and services that are based on predefined policies.

## Disaster recovery with PowerHA SystemMirror Enterprise Edition

The IBM PowerHA SystemMirror V7.1.2 Enterprise Edition helps you automate node failures and application events and provide high availability. It also helps you automate recovery actions on storage failures, controlling storage replication between sites and enabling recoveries for entire site failures, thus ensuring that copies are in a consistent state to make the failover. This enables you to build a disaster recovery solution.

The IBM PowerHA SystemMirror V7.1.2 Enterprise Edition solution can reliably orchestrate the acquisition and release of cluster resources from one site to another site. It can also provide quick failover in the event of an outage or natural disaster. PowerHA Enterprise Edition is considered a top tier disaster recovery solution because of its highly automated application failover capabilities with zero data loss and fast recovery time.

For two-site failover, there are two PowerHA inter-site cluster types, depending on the distance between the sites:

- ▶ Stretched clusters: Metropolitan area clusters that have a common SAN fabric, which allows you to share access to the storage devices from both sites.
- ▶ Linked clusters: Geographically dispersed, extended distance clusters that do not have a common SAN fabric, and the sites are connected through a WAN.

Stretched clusters typically use disk-based Metro Mirror. Being inherently synchronous, these solutions experience minimal to zero data loss, which is similar to the situation in a local cluster sharing LUNs from the same storage subsystem. Even asynchronous replication such as Global Mirror technologies can be used in this kind of stretched SAN structure if performance with synchronous replication is too slow for business requirements. But data replication cannot be ensured between sites because the I/O is completed to the application before the replication is complete.

Typically, PowerHA stretched cluster-based disaster recovery configurations for PureFlex System use WDM inter-site connectivity. Figure 8 shows the solution.

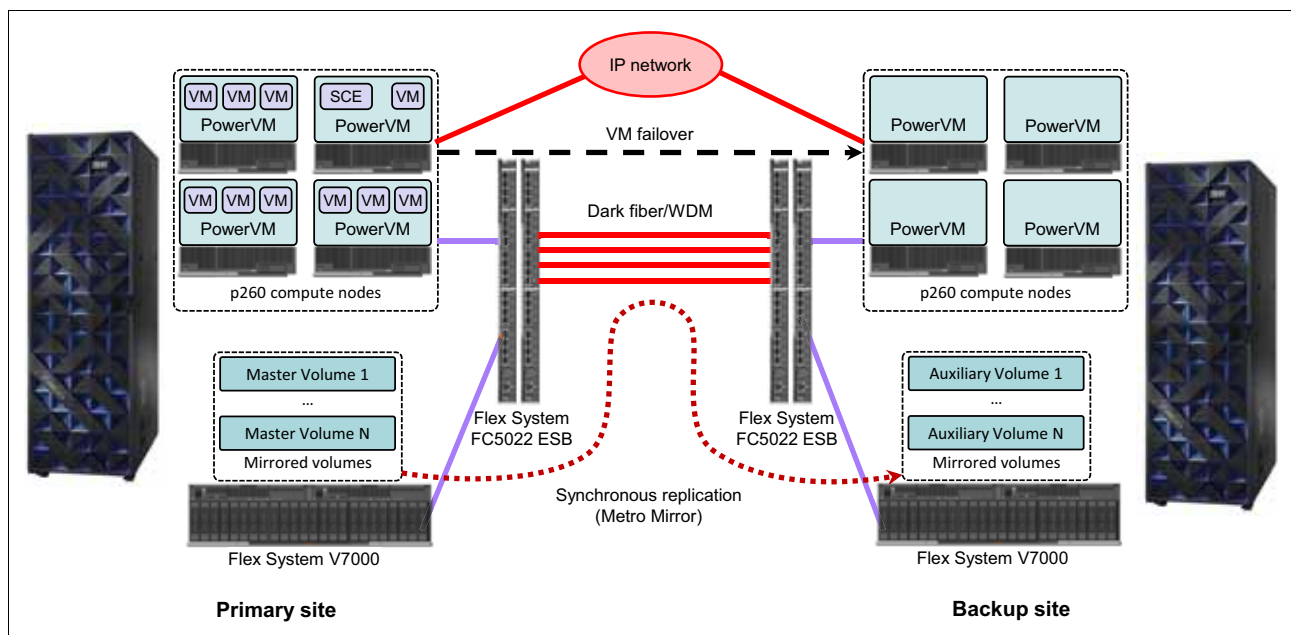


Figure 8 PowerHA SystemMirror stretched cluster

Linked clusters typically use asynchronous disk-based mirroring (Global Mirror) with FCIP inter-site connectivity, as shown in Figure 9.

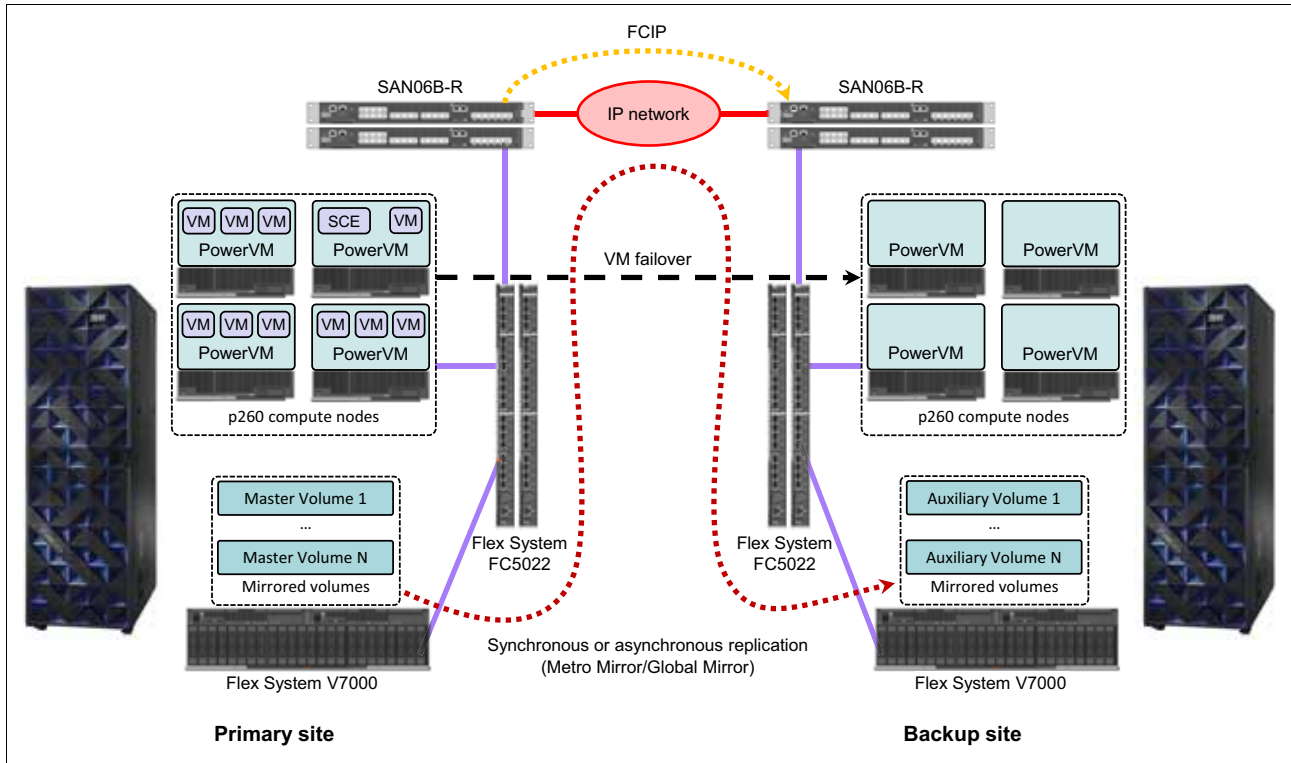


Figure 9 PowerHA SystemMirror linked cluster

In case of a disaster at the primary site, PowerHA SystemMirror performs the automated failover of applications to the backup site.

## Conclusion

IBM PureFlex System, an expert integrated system, offers a highly resilient and cloud-enabled integrated server, storage, and networking platform that supports 24x7 local operational environments, with the ability to quickly recover IT applications and data in case of disaster when combined with the Flex System V7000 replication capabilities.

The Replication Family Services of Flex System V7000 Storage Node, combined with the Extended Distance and ISL Trunking features of the Flex System FC5022 16Gb SAN switches, provide an efficient, flexible, and secure way to create remote copies of data for backup and disaster recovery purposes.

For metropolitan area distances, the usage of Metro Mirror over dark fiber/WDM or FCIP provides synchronous mirroring capabilities to support highly automated business continuity solutions with near zero RTO. For greater distances, Global Mirror provides asynchronous mirroring over FCIP to support backup/rapid restore type of disaster recovery.

Core storage management offerings, such as Tivoli Storage Manager software, can help automate backup and disaster recovery procedures for the PureFlex System cloud infrastructure. Platform-specific integration offerings, such as VMware Site Recovery Manager or PowerHA SystemMirror Enterprise Edition, can help automate site-to-site application failover in advanced disaster recovery and business continuity scenarios.



## Related publications

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

- ▶ *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844  
<http://www.redbooks.ibm.com/abstracts/sg246844.html>
- ▶ *IBM Flex System FC5022 16Gb SAN Scalable Switches*, TIPS0870  
<http://www.redbooks.ibm.com/abstracts/tips0870.html>
- ▶ *IBM PowerHA SystemMirror 7.1.2 Enterprise Edition for AIX*, SG24-8106  
<http://www.redbooks.ibm.com/abstracts/sg248106.html>
- ▶ *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984  
<http://www.redbooks.ibm.com/abstracts/sg247984.html>
- ▶ *IBM SAN Volume Controller and Brocade Disaster Recovery Solutions for VMware*, REDP-4626  
<http://www.redbooks.ibm.com/abstracts/redp4626.html>
- ▶ *IBM System Storage Business Continuity Solutions Overview*, REDP-4516  
<http://www.redbooks.ibm.com/abstracts/redp4516.html>
- ▶ *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574  
<http://www.redbooks.ibm.com/abstracts/sg247574.html>
- ▶ *Implementing SmartCloud Entry on IBM PureFlex System*, SG24-8102  
<http://www.redbooks.ibm.com/abstracts/sg248102.html>

## Acknowledgments

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

Kevin Barnes, Tamikia Barrow, Mary Comianos, Shari Deiana, Cheryl Gera, Linda Robinson, Wade Wallace, David Watts

**International Technical Support Organization, Raleigh Center**

Timothy Alpers, Bruce Corregan, Ajay Dholakia, Marco Rengan, Johann Weiss  
**IBM**

Brian Larsen, Bill Stelling, Tim Werts  
**Brocade**

## 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](http://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 IBM 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 2013. 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-5024-00 was created or updated on June 4, 2013.



Send us your comments in one of the following ways:


- ▶ Use the online **Contact us** review Redbooks form found at:  
[ibm.com/redbooks](http://ibm.com/redbooks)
- ▶ Send your comments in an email to:  
[redbooks@us.ibm.com](mailto: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](http://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®	Informix®	Redbooks (logo)  ®
BladeCenter®	Lotus®	Storwize®
DB2®	POWER®	System Storage®
Domino®	POWER7®	System x®
FlashCopy®	PowerHA®	SystemMirror®
IBM®	PowerVM®	Tivoli®
IBM Flex System™	PureFlex™	Tivoli Storage Manager FastBack®
IBM Flex System Manager™	Redbooks®	
IBM SmartCloud®	Redpaper™	

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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.