Managing IBM z/OS HyperSwap from IBM Tivoli Storage Productivity Center for Replication for Open Systems

IBM Redbooks Solution Guide

IBM® Tivoli® Storage Productivity Center for Replication is a central management solution that helps you manage replication of your critical data. It can manage hundreds of replication sessions across thousands of volumes, for both Open Systems and IBM System z® attached volumes. You can manage point-in-time and continuous volume replication services, monitor sessions, and report on volumes that are not in synchronization with the source site. Tivoli Storage Productivity Center for Replication also adds data protection capabilities with proactive management of copy services operations.

Tivoli Storage Productivity Center for Replication features include the following ones:

- Support for Metro Mirror and Global Mirror configurations and three-site recovery management
- Alerting, analysis, and reporting for improved replication management
- Central control with a simpler interface to more easily configure source-to-target pairing, run operations, and monitor progress.

Figure 1 shows the Tivoli Storage Productivity Center for Replication GUI.

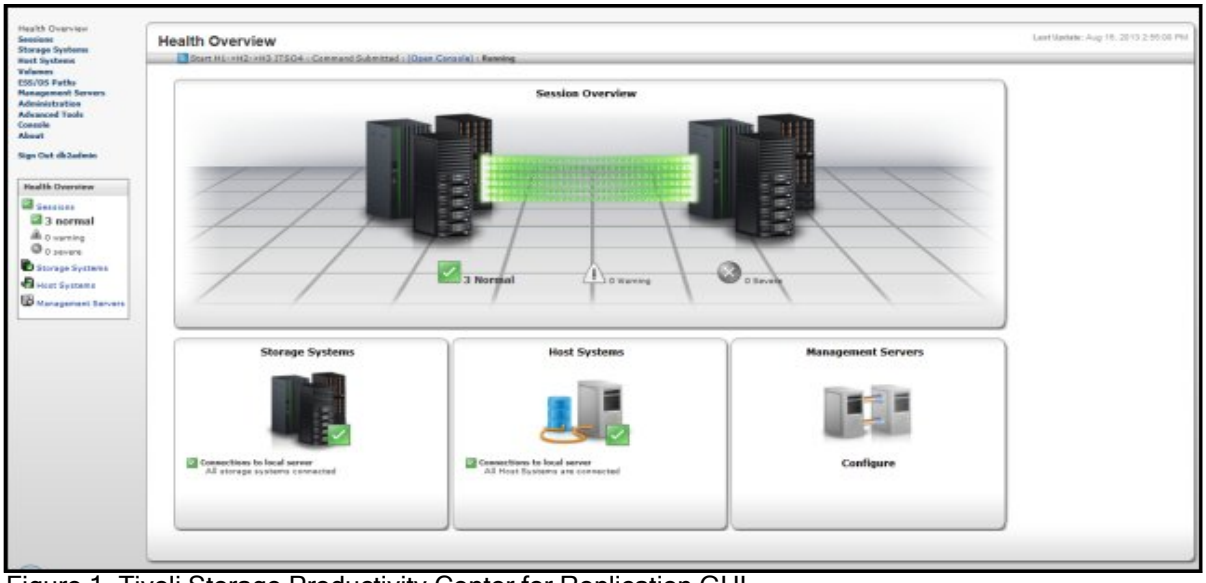

Figure 1. Tivoli Storage Productivity Center for Replication GUI

Did you know?

Tivoli Storage Productivity Center for Replication V5.2 can add a connection to an IBM z/OS® system from a Tivoli Storage Productivity Center for Replication Open Systems installation on Windows, IBM AIX®, or Linux, and manage the z/OS Basic HyperSwap® function.
On older versions of Tivoli Storage Productivity Center for Replication running on a z/OS system, you could manage only one Basic HyperSwap session from one Tivoli Storage Productivity Center for Replication server. With the new Tivoli Storage Productivity Center for Replication V5.2 running on Open Systems, you can now manage many Basic HyperSwap sessions in multiple sysplexes and monoplexes from just one Tivoli Storage Productivity Center for Replication server.

Business value

Many z/OS customers today require their applications to be available 24x7. Customers have implemented IBM Parallel Sysplex® to provide the highest availability that is available today. IBM extended the availability features of Sysplex and Parallel Sysplex to data. Storage systems today manage more data (about 50+ TB within each storage system on average). Therefore, the impact of a storage system outage is more widespread, typically impacting a Sysplex and potentially causing a Sysplex-wide outage. To address these outages, IBM developed the z/OS HyperSwap Technology, which can mask storage system failures.

HyperSwap technology is the industry standard and is provided in several different implementation options to meet the various business needs of the System z and z/OS customer base. It extends the Parallel Sysplex redundancy to storage systems. Whether your business requirements are high availability, disaster recovery, or business continuity, HyperSwap technology provides a solution for your environment.

The HyperSwap technology provides planned and unplanned actions in a z/OS environment.

Planned HyperSwap actions can:

- Transparently switch all primary storage systems disks with the auxiliary storage systems disks
- Perform storage systems configuration maintenance and planned site maintenance without requiring any applications to be quiesced

Unplanned HyperSwap actions:

- Can transparently switch to auxiliary storage systems if there are unplanned outages of the primary storage systems.
- Allow production systems to remain active during a storage system failure. The storage system failures no longer constitute a single point of failure for an entire Parallel Sysplex.

Tivoli Storage Productivity Center for Replication V5.2 enables you to manage z/OS Basic HyperSwap and helps you manage planned and unplanned actions in an z/OS environment from an Open Systems environment. It manages the z/OS Basic HyperSwap function with code in the z/OS systems, that is, the Input/Output Supervisor (IOS) component.

IBM has done an analysis of field storage system failures, and as a result created a set of trigger events that are monitored by Tivoli Storage Productivity Center for Replication. When one of these HyperSwap trigger events occurs, a *data freeze* across all logical subsystems (LSSs) on all storage systems is started. All I/Os to all devices are queued (extended long busy state), thus maintaining full data integrity and cross volumes data consistency. z/OS then completes the Basic HyperSwap function of recovering the target devices and rebuilding all z/OS internal control blocks to point to the recovered target devices. When this task completes, all I/O is released and all applications continue to run against the recovered target devices, thus transparently managing a complete storage system outage, with a dynamic *busy* state and a redirection of all host I/O. Applications must be able to tolerate the extended long busy state, which is transparent to the applications, but elongates I/O that is in progress, until the HyperSwap actions are complete.

Solution overview

HyperSwap technology can be used in high availability, disaster recovery, or business continuity solutions with two or three sites, and is based on synchronous replication. By definition, synchronous replication makes the issuing host wait for its write I/Os to complete at both the local and remote storage systems. The solutions are managed by Tivoli Storage Productivity Center for Replication servers in a high availability configuration that manages z/OS Basic HyperSwap on z/OS systems in a Sysplex environment.

The following Tivoli Storage Productivity Center for Replication sessions use HyperSwap technology:

- Basic HyperSwap
- Metro Mirror with Failover/Failback
- Metro Global Mirror
- Metro Global Mirror with Practice

The overview of the solution that is based on Tivoli Storage Productivity Center for Replication is shown in Figure 2.

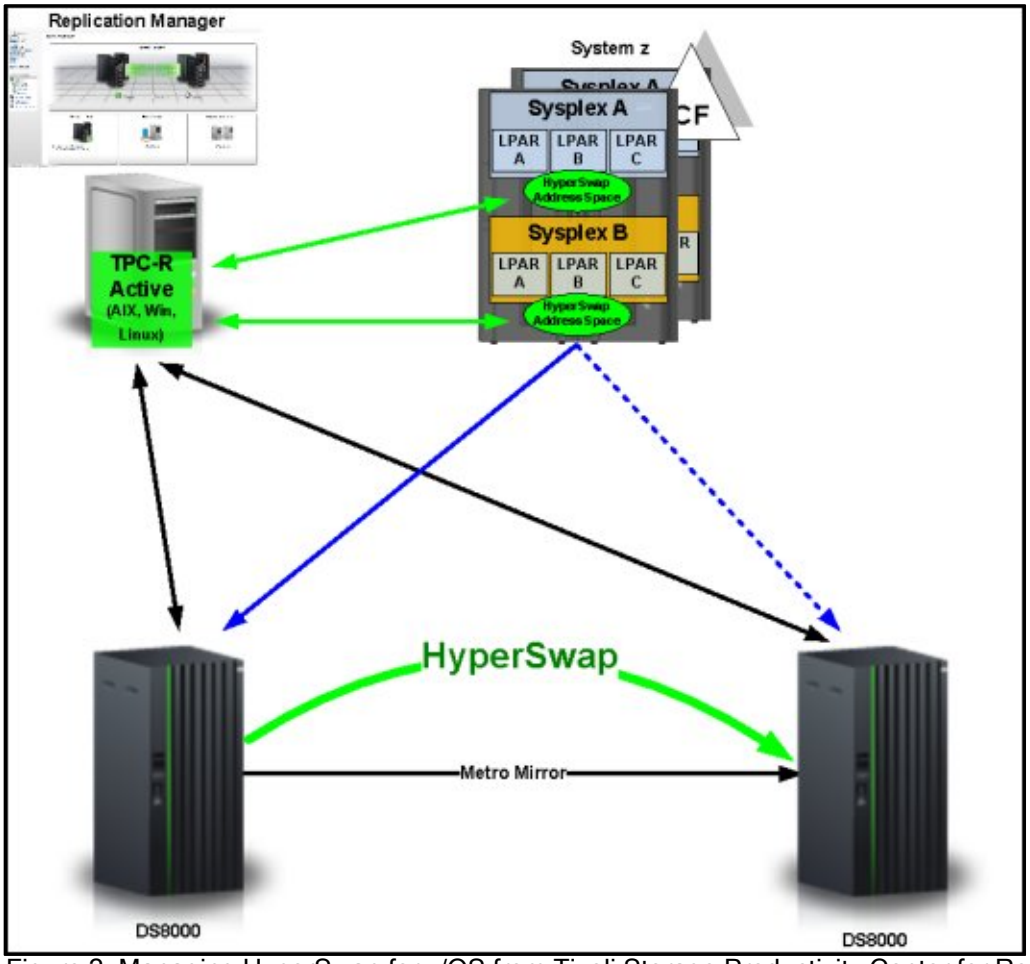

Figure 2. Managing HyperSwap for z/OS from Tivoli Storage Productivity Center for Replications Open Systems

Solution architecture

In this solution, an Active-Standby Tivoli Storage Productivity Center for Replication configuration is deployed across the two sites, with the active server running on the primary storage site (Site A). An IP connection is provided to manage the storage systems and to communicate with the z/OS HyperSwap address spaces. A Metro Mirror Failover/Failback session with HyperSwap enabled is defined to Tivoli Storage Productivity Center for Replication to manage this kind of configuration. Along with the Parallel Sysplex facilities, this Tivoli Storage Productivity Center for Replication implementation can provide high availability features that cover many failure scenarios.

Figure 3 shows a possible implementation of Tivoli Storage Productivity Center for Replication for Open Systems to manage data availability.

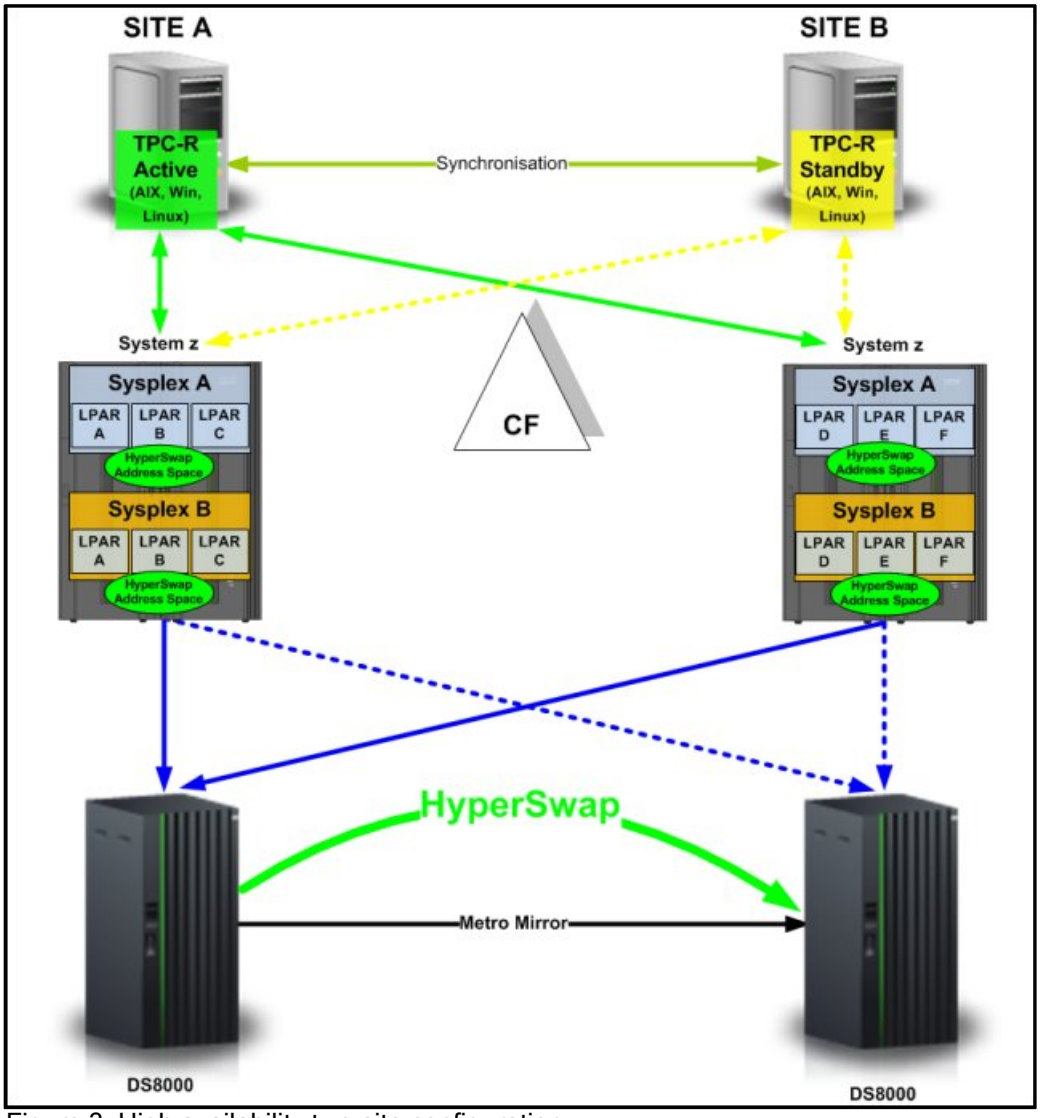

Figure 3. High availability two-site configuration

Tivoli Storage Productivity Center for Replication can also be implemented for Open Systems to manage the data availability in a three-site solution, as shown in Figure 4. In this solution, a high availability Tivoli Storage Productivity Center for Replication configuration is deployed across two sites, with the active server running on the primary storage site (Site A) and the standby server running at the remote site (Site C). An IP connection is provided to manage the storage systems and to communicate with the z/OS HyperSwap address spaces. A Metro Global Mirror session with HyperSwap enabled is defined to Tivoli Storage Productivity Center for Replication to manage this configuration. This Tivoli Storage Productivity Center for Replication implementation, along with planned events, provides protection for unplanned events.

Figure 4 illustrates a possible solution architecture that manages data availability by using three sites.

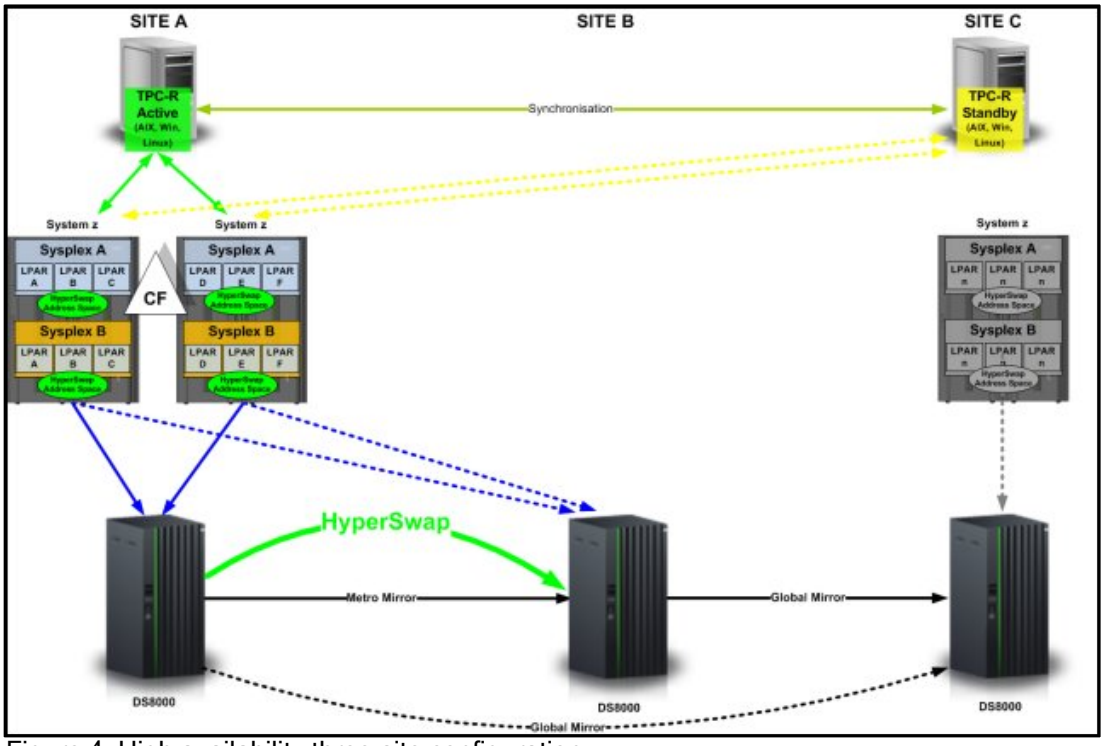

Figure 4. High availability three-site configuration

Usage scenarios

This section describes two usage scenarios for this solution.

Scenario 1

An active-active campus is a high availability solution that is based on advanced clustering and data replication technologies. This scenario provides continuous availability features by implementing hardware redundancy in multiple locations and sophisticated software facilities that allow workload switching among the sites. Also, the data replication plays an important role in this solution by providing a consistent copy of data across the campus.

In a typical active-active campus configuration, the application workload is running in two separate sites that are interconnected through LAN and SAN extensions. Consistent copies of the data are continuously present in both sites. This solution can manage, transparently to the users, various critical situations such as single or multiple hardware failures, or even an entire site failure. In terms of business continuity, the active-active campus configuration cannot be considered a disaster recovery solution because the connection latency impact requires a relatively short distance between the campus sites, and both sites might not survive the same major natural disasters.

Figure 5 shows a schematic representation of an active-active campus configuration that is based on Tivoli Storage Productivity Center for Replication, IBM System z, and IBM System Storage® DS8000® solutions.

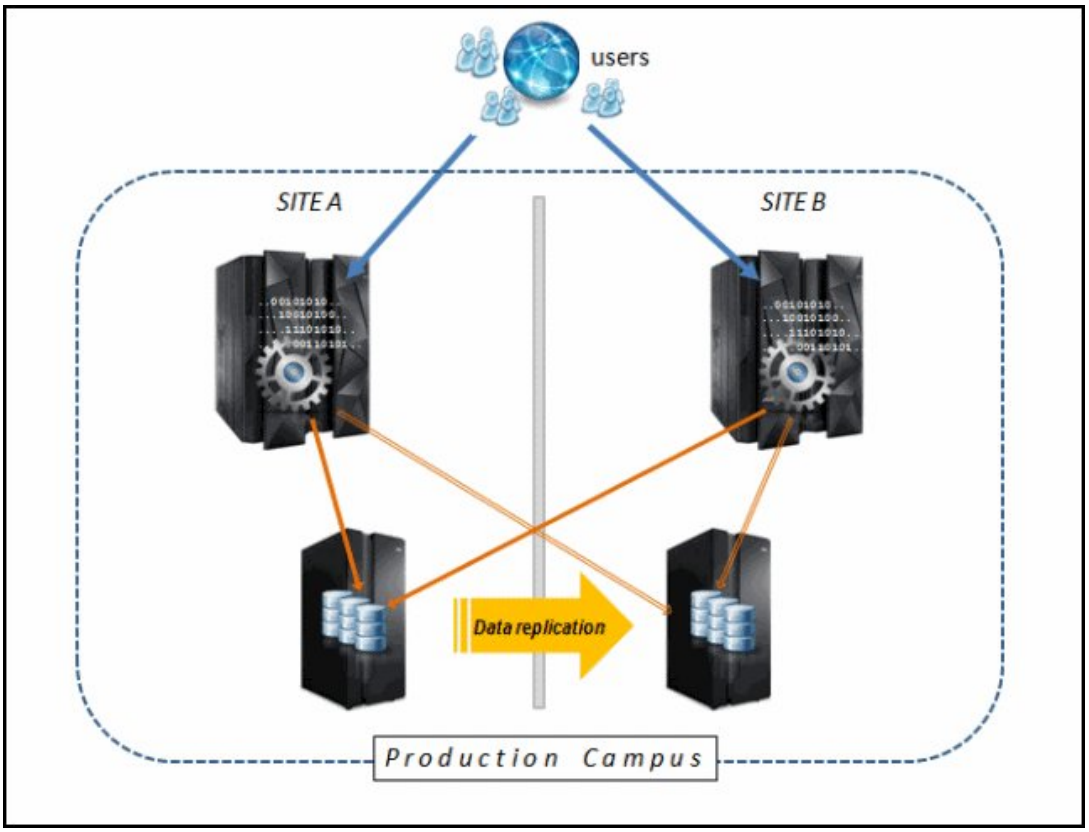

Figure 5. Active-active campus configuration that is based on Tivoli Storage Productivity Center for Replication, System z, and DS8000 solutions

In this scenario, the application workload is running at both Site A and Site B. The Parallel Sysplex facility manages the availability of the application across the sites, and the data is replicated using synchronous mirroring technologies such as DS8000 Metro Mirror. To manage the data replication and the data availability, a Tivoli Storage Productivity Center for Replication for Open Systems solution can be implemented.

Figure 6 shows a primary storage system failure at Site A. In this case, the failure of a primary storage system triggers a HyperSwap switch that allows application transparent switching to the auxiliary storage system at Site B.

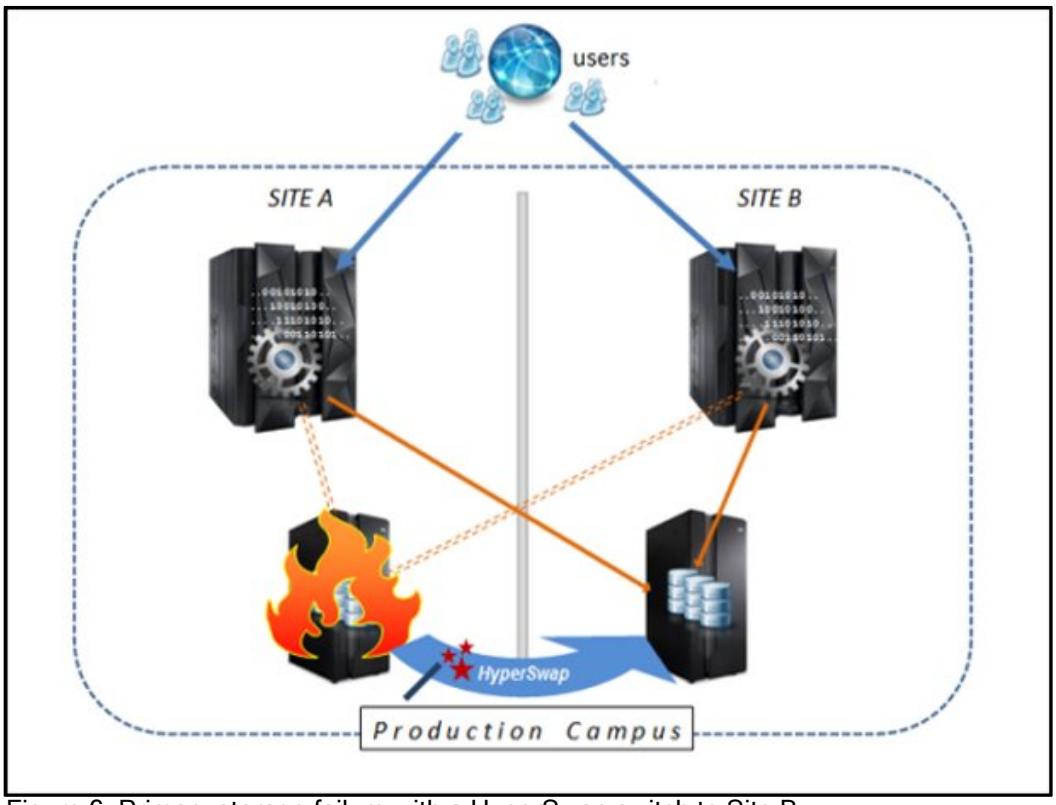

Figure 6. Primary storage failure with a HyperSwap switch to Site B

Figure 7 shows a complete Site A failure. Again, the HyperSwap capabilities of Tivoli Storage Productivity Center for Replication let the applications running on Site B survive the Site A failure by switching the I/O to the Site B storage system. In addition, the Parallel Sysplex facilities allow application switching from Site A to Site B with minimal or even no downtime.

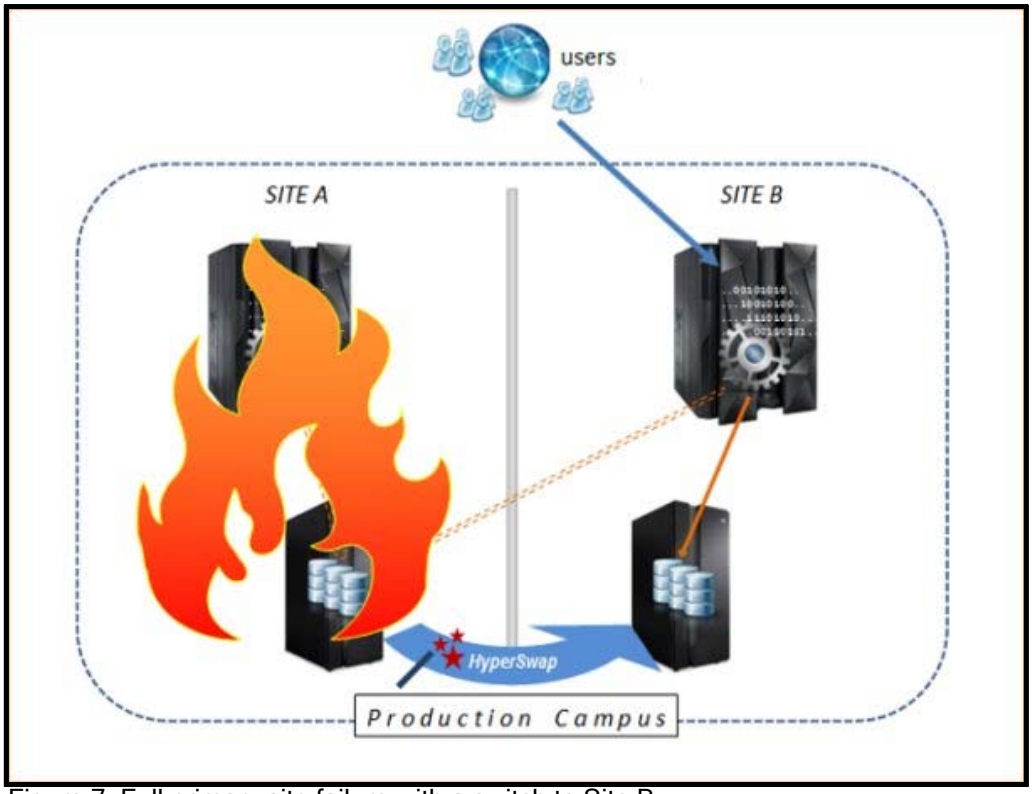

Figure 7. Full primary site failure with a switch to Site B

Scenario 2

A three-site configuration, which is shown in Figure 8, combines the high availability features of an active-active campus configuration with disaster recovery capabilities. In addition to the infrastructure that is required for the campus solution, a remote third site is set up with computing capabilities and storage capacity for disaster recovery purposes. To make this solution effective in terms of Recovery Time Objective (RTO) and Recovery Point Objective (RPO), additional data replication capabilities must be implemented to ensure data availability at the remote site. Furthermore, to be considered effective in terms of disaster recovery, a three-site solution must have enough distance between the campus sites and the remote site such to ensure the survival of at least one site in case of a major disaster. For this reason, the distance between the campus sites and the remote site does not allow the use of synchronous replication technologies because of the latency impacts on write operations.

In the scenario that is shown in Figure 8, the application workload is running in the campus infrastructure. Asynchronous mirroring technologies, such as DS8000 Global Mirror, are used to replicate the data from the intermediate site (Site B) to the remote site (Site C). Combining HyperSwap capability and Global Mirror management facilities, Tivoli Storage Productivity Center for Replication for Open Systems provides a single point of control for data availability management, monitoring, and alerting.

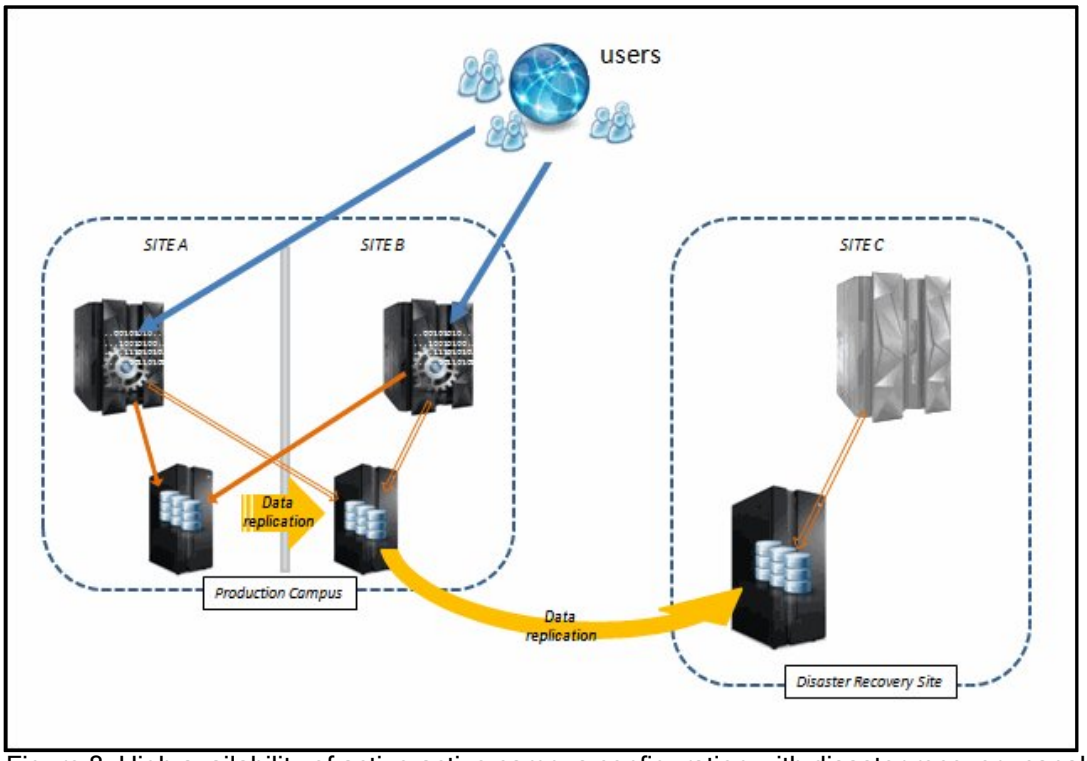

Figure 8. High availability of active-active campus configuration with disaster recovery capabilities

Figure 9 a complete failure of Site B. In this case, the disaster affects only the application that is running in Site B, and the Parallel Sysplex facilities can minimize the impact to the business. From the disaster recovery point of view, Tivoli Storage Productivity Center for Replication can restart the replication to the remote site directly from the primary site without requiring a full copy of the data by using the Global Mirror Incremental Resync feature.

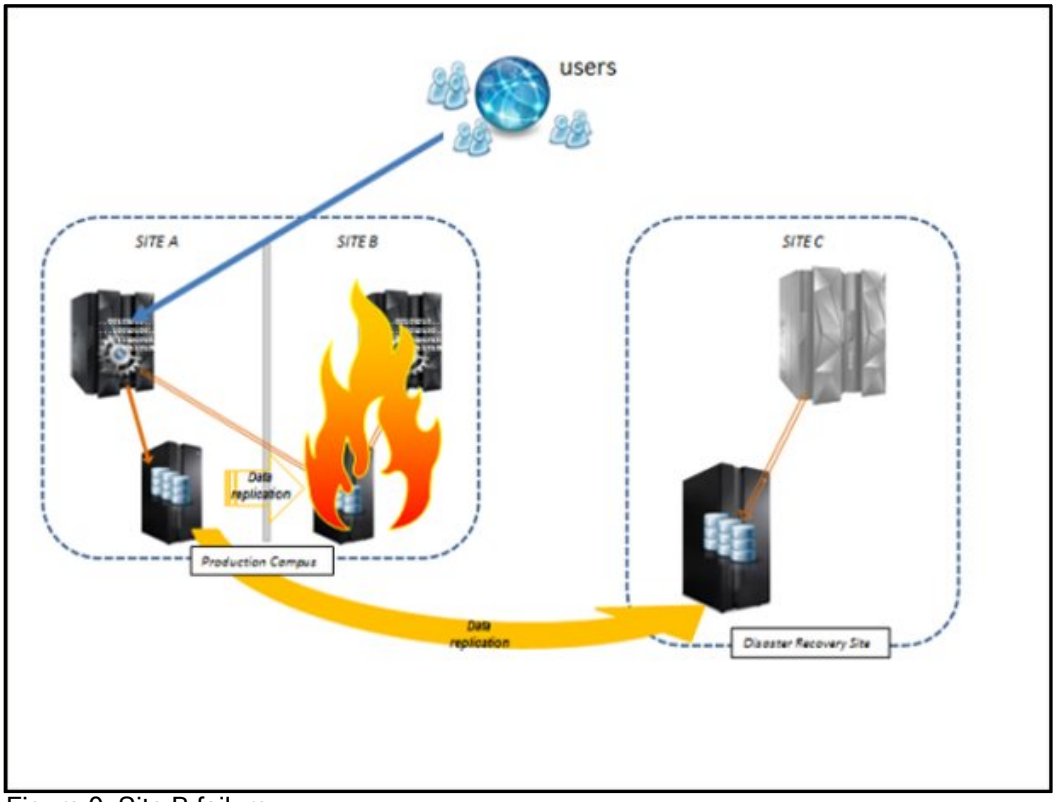

Figure 9. Site B failure

Figure 10 shows a complete production campus failure scenario. A complete production campus failure requires a disaster recovery scenario that can perform a full recovery of the operations at the remote site. Tivoli Storage Productivity Center for Replication offers all the capabilities that are needed to perform the data recovery operations, and all the functions to go back to the original three-site configuration (Go-home procedures).

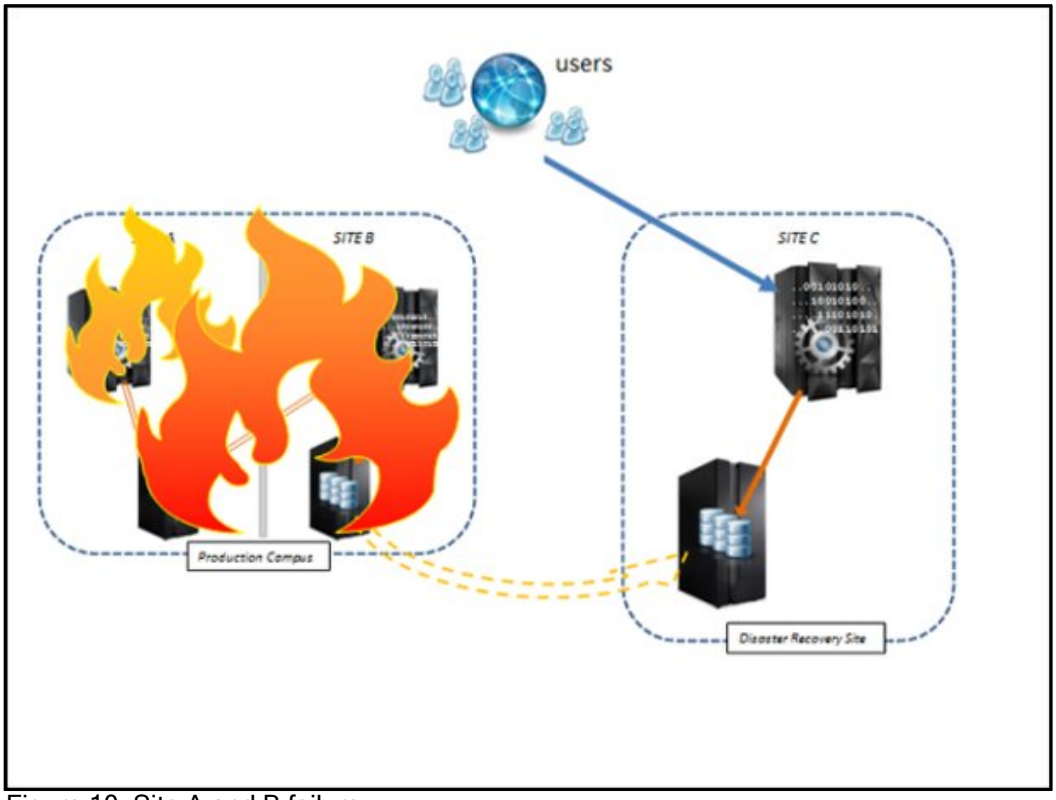

Figure 10. Site A and B failure

Integration

Tivoli Storage Productivity Center is an industry-leading storage resource management software that provides comprehensive visibility, control, and automation for managing heterogeneous storage infrastructures through a centralized, web-based management console. It provides storage infrastructure management capabilities such as automated system discovery, provisioning, configuration management, performance monitoring, and replication for storage systems and storage networks. Tivoli Storage Productivity Center provides storage administrators a simple way to conduct device management for multiple storage arrays and SAN fabric components.

Tivoli Storage Productivity Center

Tivoli Storage Productivity Center for Replication is a component of the Tivoli Storage Productivity Center that is responsible for managing the replication of storage systems. Tivoli Storage Productivity Center monitors replication and it shows different replication alerts that are triggered by specific conditions. This integration of Tivoli Storage Productivity Center and Tivoli Storage Productivity Center for Replication provides a complete storage resource management solution.

IBM SmartCloud Virtual Storage Center

IBM SmartCloud® Virtual Storage Center software is another solution that integrates Tivoli Storage Productivity Center for Replication. IBM SmartCloud Virtual Storage Center includes core functions from three top IBM offerings:

- Advanced functions from IBM Tivoli Storage Productivity Center, where the Tivoli Storage Productivity Center for Replication is a component
- External virtualization, which is found in IBM System Storage SAN Volume Controller
- Application-aware snapshot backup and restore capabilities from IBM Tivoli Storage FlashCopy® Manager

IBM SmartCloud Virtual Storage Center enables all three smarter storage characteristics (virtualization, advanced GUI, and near-instant and application-aware backup and restore) for existing storage infrastructures. Its built-in efficiency features help users avoid the need to purchase add-ons or additional licenses, or deal with complicated integration issues.

Supported platforms

Tivoli Storage Productivity Center for Replication and HyperSwap sessions require IBM Storage System DS8000, IBM System Storage DS6000™, or IBM System Storage ESS 800 with the Metro Mirror advanced copy function.

z/OS HyperSwap support requires z/OS V1.12 or later, and it requires that HyperSwap address spaces, the HyperSwap Management address space, and the HyperSwap API address space are running. These address spaces must be running in z/OS systems and sysplexes that use the HyperSwap function. The Basic HyperSwap function is supported

A detailed list product and platform support for Tivoli Storage Productivity Center for Replication V5.2 can be found at the following links:

- IBM Support web page:

  http://www-01.ibm.com/support/docview.wss?uid=swg21386446

- IBM Tivoli Storage Productivity Center information center link:

  http://pic.dhe.ibm.com/infocenter/tivihelp/v59r1/index.jsp

Ordering information

This section describes the licensing and ordering information for IBM Tivoli Storage Productivity Center and IBM SmartCloud Virtual Storage Center. Table 1 provides the ordering part numbers and feature codes.

IBM SmartCloud Virtual Storage Center

License function title: This program is licensed under the IBM Program License Agreement (IPLA) and the associated Agreement for Acquisition of Software Maintenance, License Information Form number L-JSHW-8XNF7U, available for review on the IBM Software License Agreement website:

http://www.ibm.com/software/sla/sladb.nsf

Product group: Tivoli Storage
Product category: Tivoli Storage Productivity Center

Ordering information is shown in Table 1.

Table 1. Ordering part numbers and feature codes

Program name	Version	PID number
IBM Tivoli Storage Productivity Center	5.2	5725-F93
IBM SmartCloud Virtual Storage Center (for IBM Passport Advantage®)	5.2	5725-F92
IBM SmartCloud Virtual Storage Center (For Advanced Administrative System (AAS))	5.2	5608-AE1

Related information

For more information, see the following documents:

- IBM Offering Information page (to search on announcement letters, sales manuals, or both): http://www.ibm.com/common/ssi

  On this page, enter the solution name of interest (SmartCloud Virtual Storage Center, Tivoli Storage Productivity Center, and so on), select the information type, and click **Search**. On the next page, narrow your search results by geography and language.

- *IBM Tivoli Storage Productivity Center for Replication for Series z*, SG24-7563 (currently in draft form only): http://www.redbooks.ibm.com/redpieces/abstracts/sg247563.html?Open
- *IBM Tivoli Storage Productivity Center for Replication for Open Systems*, SG24-8149 http://www.redbooks.ibm.com/redpieces/abstracts/sg248149.html?Open
- Supported Products and Platforms Interoperability Matrix Links: http://www-01.ibm.com/support/docview.wss?uid=swg21386446

Notices

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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 measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

COPYRIGHT LICENSE:

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

This document was created or updated on December 18, 2013.

Send us your comments in one of the following ways:

- Use the online **Contact us** review form found at:
  ibm.com/redbooks
- Send your comments in an e-mail to:
  redbook@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.

This document is available online at http://www.ibm.com/redbooks/abstracts/tips1068.html .

Trademarks