Mirror, mirror (on the site)

Deployment scenarios for TPC-R on z/OS

BY NICK CLAYTON

Introduction

TotalStorage® Productivity Center for Replication (TPC-R) is the IBM™ enterprise replication management solution. TPC-R can manage a range of functions for the enterprise disk platform (DS8000™, DS6000™ and ESS) as well as the San Volume Controller (SVC). This includes synchronous and asynchronous replication, point-in-time copy function, and combinations such as 3-site solutions and continuous replication with an integrated point-in-time copy for testing purposes.

TPC-R is a WebSphere application that can run on a variety of platforms including z/OS®, where it runs as a z/OS UNIX® System Services (USS) application. From an operational perspective, it provides the same function as that on any other platform but also has a number of capabilities and considerations unique to the z/OS platform that this article will discuss in more detail.

How TPC-R works

TPC-R uses the concept of sessions of various types to define the set of devices to be replicated and to help simplify the operation of the underlying replication. TPC-R provides a graphical user interface as well as a command line interface for integration with automation and scheduled operations. Alerting capabilities are also provided using Simple Network Management Protocol (SNMP). For more information, see *IBM TotalStorage Productivity Center for Replication for Series z* at the following Web site:
http://www.redbooks.ibm.com/redpieces/abstracts/sg247563.html?Open

Metro Mirror and TPC-R

Metro Mirror provides a synchronous replication capability for IBM enterprise disk subsystems and a solution for both high availability and disaster recovery. One of the key components of the Metro Mirror architecture is the consistency group function. This ensures that if there is a suspension of the Metro Mirror device pair, the whole environment is suspended such that all the secondary devices are consistent with each other. This is supported for both planned and unplanned situations.

When running TPC-R on z/OS with a Metro Mirror session, there are some special considerations regarding this consistency group function. In a distributed systems environment, it is unlikely that the server on which TPC-R is running has its data replicated by that TPC-R server; however, in a z/OS environment, this is more likely the case, as it would be usual to run the TPC-R application on an existing z/OS system.

If you are using a Metro Mirror session to replicate data, it is essential that the TPC-R server is not running on a z/OS image where Metro Mirror is “mirroring” either the system or TPC-R server data. Such a configuration prevents the TPC-R server from correctly completing the consistency group automation function, which would be impacted by the suspension and the consistency group creation event known as a *freeze.* This freeze can result in up to a two-minute *long busy* period for the production volumes.

One of the other features supported by TPC-R is the capability to define a backup server that you can activate in case of an outage of the primary server. It is possible to run this server on a different platform from the primary server to allow additional flexibility in deploying a TPC-R environment. Even if the primary server is running on z/OS, it might be desirable to run the backup server on a distributed platform to reduce costs and the overhead of having an active z/OS image in the recovery location.

Figure 1 shows how the setup of a Metro Mirror session with TCP-R should look:

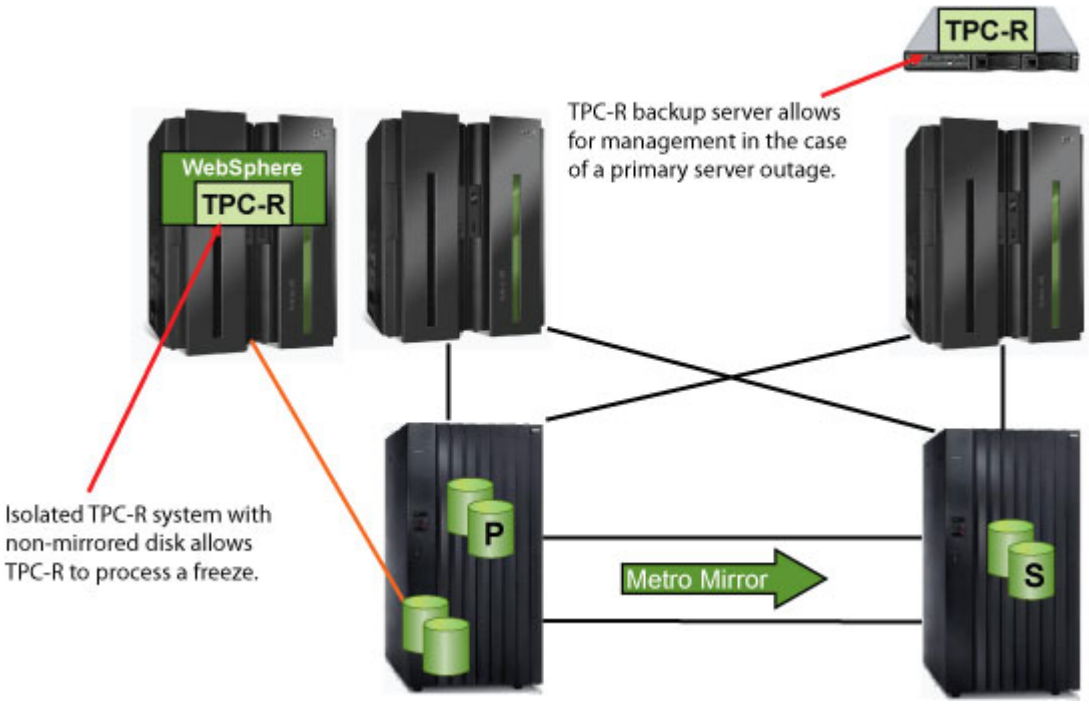

Figure 1: Metro Mirror environment with TPC-R on z/OS

Basic HyperSwap and TCP-R

While disk replication solutions such as Metro Mirror can provide a high availability solution, they do not provide continuous availability in case of disk subsystem failures. In this case, the system must be shut down and then restarted on the secondary devices. Basic HyperSwap™ helps solve this problem by providing a solution for continuous availability for z/OS data within a single site. Basic HyperSwap can transparently switch the workload of a sysplex from the primary volumes of a Metro Mirror environment to continue to run on the secondary volumes.

Basic HyperSwap function can run with the standard TPC-R but also works with an entitled version called TPC-R Basic Edition. Note that this version only supports the Basic HyperSwap session type and does not include other TPC-R functions such as the ability to define a backup TPC-R server.

Components of Basic HyperSwap

The Basic HyperSwap solution has three main components:

- TPC-R server, which contains the configuration and is used for management of the environment. This component must run on a z/OS image that must be a member of the sysplex using Basic HyperSwap.
- One or more Basic HyperSwap address spaces, which control and coordinate the HyperSwap function and run on all members of the sysplex
- One or more HyperSwap API Address spaces, which execute the various actions required for HyperSwap processing and run on all members of the sysplex.

See Figure 2:

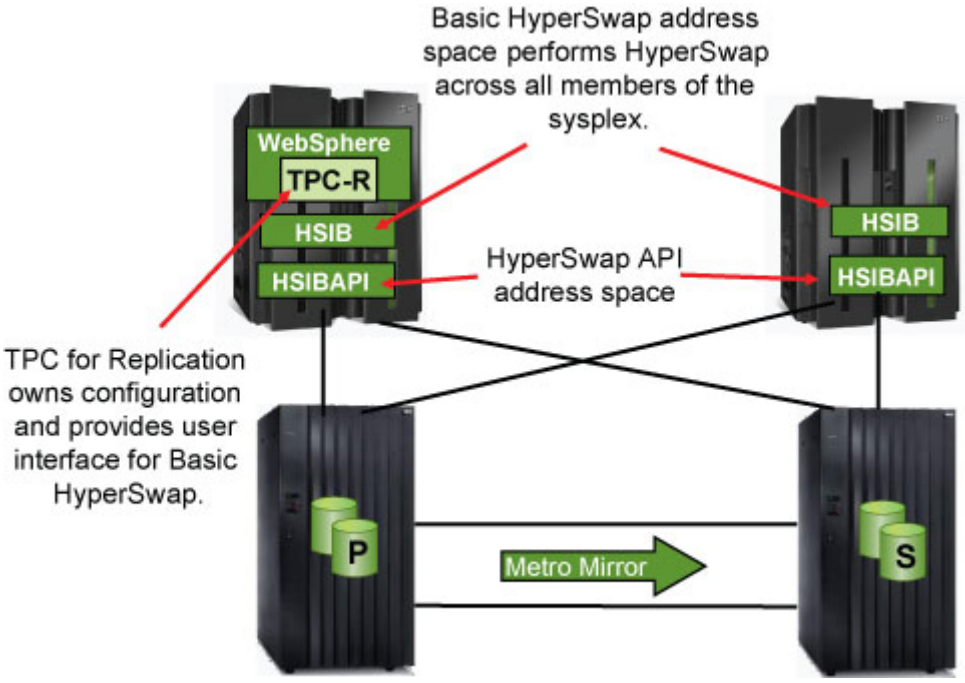

Figure 2: Basic HyperSwap environment with TPC-R on z/OS

Basic HyperSwap is not designed to provide a disaster recovery capability; thus, if there is replication failure, the consistency group function is not used to consistently suspend the secondary devices. The HyperSwap address space ensures that the production I/O is allowed to continue on the primary devices, but it does not cause a consistent suspend.

For more information about Basic HyperSwap, see the Redpaper *IBM DS8000 and z/OS Basic HyperSwap* at the following Web site:
http://www.redbooks.ibm.com/abstracts/redp4441.html?Open

Basic HyperSwap with Metro Global Mirror

Metro Global Mirror is a cascaded mirroring solution that uses a combination of the Metro Mirror synchronous replication capability and the asynchronous replication provided by Global Mirror. This duality provides both a local synchronous solution for high availability and a long-distance solution to protect against regional events that might impact both synchronous copies. TPC-R V4.1 introduced the capability for TPC-R to manage a 3-site Metro Global Mirror environment and also to provide HyperSwap for the Metro Mirror relationships.

In this configuration, the Basic HyperSwap address space provides both the HyperSwap coordination and the consistency group function. This solution provides both a local high availability and disaster recovery capability with a long-distance replication function to handle regional disasters. Because the HyperSwap address space is able to operate even if all devices are in a long busy state or otherwise unavailable, you don’t need an isolated system to provide the consistency group function.

See Figure 3:

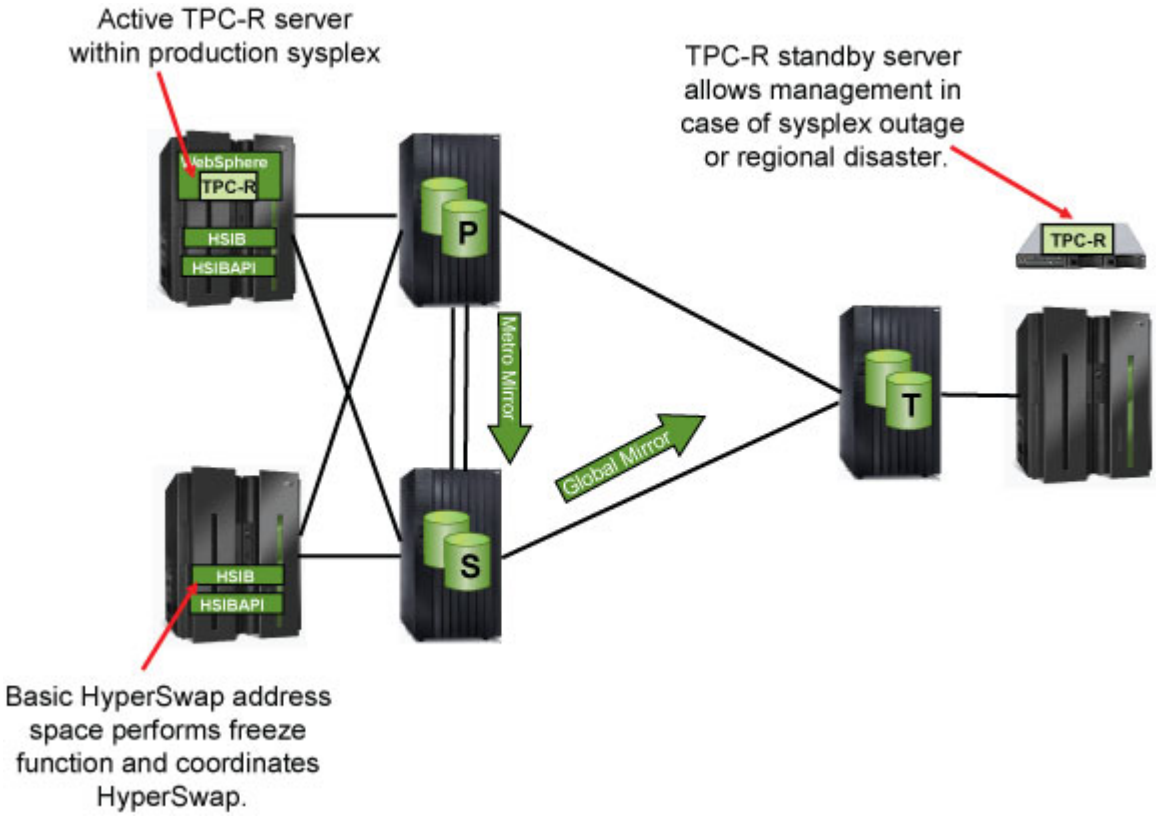

Figure 3: Metro Global Mirror environment with TPC-R on z/OS

In a Metro Global Mirror environment, the standby TPC-R server should be located in the third remote site. This configuration allows you to control the recovery of the environment if the two local sites fail. It also allows the standby server to perform the necessary actions if a primary site failure also results in a sysplex outage or failure of the primary TPC-R server system.

Conclusion

The HyperSwap function provided by TPC-R on z/OS provides significant benefits compared to running the solution on a distributed systems platform while still providing the same operational interface and the ability to manage replication on other platforms and disk subsystems.

TPC-R only provides a replication management solution, not a complete high availability and disaster recovery solution, as it does not contain function to manage the systems and restart workloads in the event of a failure. If this is required, the various GDPS® solutions can provide this capability. More information regarding GDPS can be found in *GDPS Family - An Introduction to Concepts and Capabilities* at the following Web site:
http://www.redbooks.ibm.com/abstracts/sg246374.html

Author biography:

Nick Clayton is a Consulting IT Specialist working for the European Storage Competence Centre where he specializes in high availability and disaster recovery solutions. Nick graduated from Trinity College Cambridge in 1994 with a degree in Mathematics and has been with IBM for 10 years.