IBM® Storage

IBM FlashSystem HyperSwap with SAP HANA on IBM Power

Version 1.1

IBM Storage Team

Contents

About this document

This paper is intended as an IBM® solution reference for using IBM FlashSystem® with SAP HANA® in an SAP HANA tailored data center integration (SAP HANA TDI) environment. SAP HANA TDI allows the SAP customer to attach external storage to the SAP HANA server. This paper demonstrates the storage configuration with IBM HyperSwap® and IBM FlashSystem, which provides a high availability solution for SAP HANA.

This document is written for system integrators, system or storage administrators, customers, and business partners with knowledge about SAP HANA and IBM System Storage®.

Scope

This document was developed using the following software tools:

- SAP HANA 2.0 database
- IBM Storage Insights®
- SAP HANA Hardware and Cloud Measurement Tools (HCMT)
- IBM Spectrum® Virtualize
- IBM HyperSwap

This technical report does not:

- Replace any official manuals and documents issued by IBM
- Explain installation and configuration of SAP HANA

Prerequisites

For a list of all IBM storage systems certified for SAP HANA production please refer to:

- Certified and supported SAP HANA hardware

It is assumed that you are familiar with and have basic knowledge of the following products:

- IBM FlashSystem
- SAP HANA database
- IBM Power® server

IBM Storage Insights

IBM Storage Insights is offered free of charge to customers who own IBM block storage systems. It is a secured IBM Cloud® storage service that monitors IBM block storage and provides advanced functionality for alerting.

It provides single-pane views of IBM block storage systems, such as the Operations dashboard and the Notifications dashboard. With the information that is provided, such as the diagnostic event information, key capacity and performance information, and the streamlined support experience, you can quickly assess the health of your storage environment and get help with resolving issues. Furthermore, on the Advisor page, IBM Storage Insights provides recommendations on the remedial steps that can be taken to manage risks and resolve issues that might impact your storage services.

All IBM FlashSystem performance graphs in this paper were created with IBM Storage Insights.

For more information about IBM Storage Insights see the following resources:

https://www.ibm.com/products/analytics-driven-data-management

IBM FlashSystem

The IBM FlashSystem family combines the performance of flash and end-to-end Non-Volatile Memory Express (NVMe) with the reliability and innovation of IBM FlashCore technology, the ultra-low latency of Storage Class Memory (SCM), the rich features of IBM Spectrum Virtualize and AI predictive storage management and proactive support by Storage Insights. Built in a powerful 1-2 U enterprise-class, blazing fast storage all-flash array, as shown in Figure 1.

NVMe protocol inside FlashSystem
NVM Express (NVMe) is an optimized, high-performance scalable host controller interface designed to address the needs of systems that utilize PCI Express-based solid-state storage. The NVMe protocol is an interface specification for communicating with storage devices. It is functionally analogous to other protocols, such as SAS. However, the NVMe interface was designed for extremely fast storage media, such as flash-based solid-state drives (SSDs) and low-latency non-volatile storage technologies.

NVMe storage devices are typically directly attached to a host system over a PCI Express

(PCIe) bus. That is, the NVMe controller is contained in the storage device itself, alleviating the need for an additional I/O controller between the CPU and the storage device. The architecture results in lower latency, throughput scalability, and simpler system designs. NVMe protocol supports multiple I/O queues, versus legacy SAS and SATA protocols that use only a single queue.

These all-flash systems include IBM Spectrum Virtualize software and introduce remarkable new features in comparison to the predecessor models:

- End-to-end **NVMe** support: NVMe is a logical device interface standard from 2011 for accessing non-volatile storage media that is attached via a PCI Express bus.
- **Lower latencies** through RDMA: Direct memory access from the memory of one node into that of another without involving either one's operating system.
- **Data reduction pools** (DRP) represent a significant enhancement to the storage pool concept. Now with the introduction of data reduction technology, compression, and deduplication, it has become more of a requirement to have an uncomplicated way to stay "thin".
- **FlashCore Modules** (FCMs) or industry standard NVMe drives can be used for the IBM FlashSystems. If the FCM option is chosen, then the user can take advantage of the built-in hardware compression, which will automatically try to compress the stored data when written to the drives.
- Thin-provisioned **IBM FlashCopy** uses disk space only when updates are made to the source or target data, and not for the entire capacity of a volume copy.
- **HyperSwap** capability enables each volume to be presented by two IBM FlashSystems. This high-availability configuration tolerates combinations of node and site failures, using host multipathing driver, based on the one that is available for the regular IBM FlashSystem.
- The IBM FlashSystem supports the new low latency, high speed Storage Class Memory (SCM). SCM is a non-volatile memory device that performs faster (~10µs) than traditional NAND SSDs(100µs), but slower than DRAM (100ns).
- IBM Storage Insights is an additional part of the monitoring capability of the IBM FlashSystem family and supplements the views available in the GUI.

Figure 1: IBM FlashSystem 2U control enclosure (Models 9200 and 7200)

HyperSwap architecture

HyperSwap is the high availability (HA) solution for IBM FlashSystem that provides continuous data availability in case of hardware failure, power failure, connectivity failure, or disasters. HyperSwap functions are available on IBM FlashSystems that can support more than one I/O group, such as the IBM FlashSystem 5015, 5035, 5200, 7200, 9200 as shown in Figure 2. The IBM FlashSystem provides GUI and CLI management of the HyperSwap function.

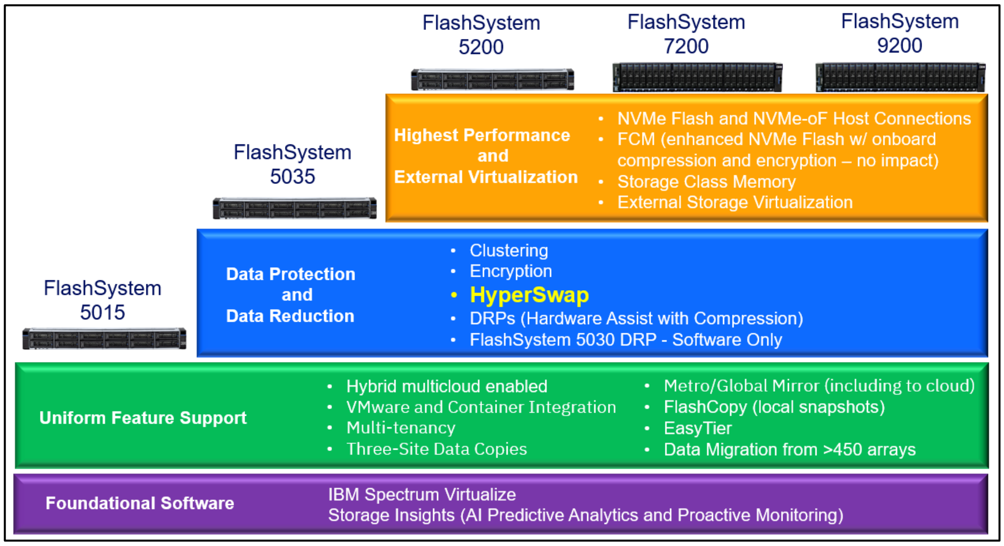

Figure 2: IBM FlashSystem family

IBM HyperSwap has the following key features:

- Works with all IBM Spectrum Virtualize products (IBM SAN Volume Controller and IBM FlashSystem) except for IBM FlashSystem 5015.
- Uses intra-cluster synchronous Remote Copy (named Active-Active Metro Mirror) capability along with change volumes and access I/O group technologies.
- Makes a host's volumes accessible across two IBM FlashSystem I/O groups in a clustered system by using the Active-Active Metro Mirror relationship. The volumes are presented as a single volume to the host.
- Works with the standard multipathing drivers that are available on various host types, with no additional host support required to access the highly available volumes.

With HyperSwap, a fully independent copy of the SAP HANA data, log and shared volumes is maintained at each site. When data is written by hosts at either site, both copies are synchronously updated before the write operation is completed. The HyperSwap function automatically optimizes itself to minimize data that is transmitted between two sites, and to minimize host read and write latency. The HyperSwap solution requires one IBM FlashSystem Control Enclosure at each site, and it requires a third site that acts as a tie-breaking quorum device. The third site can be implemented as fibre channel attached storage or IP-linked quorum application, as shown in Figure 3.

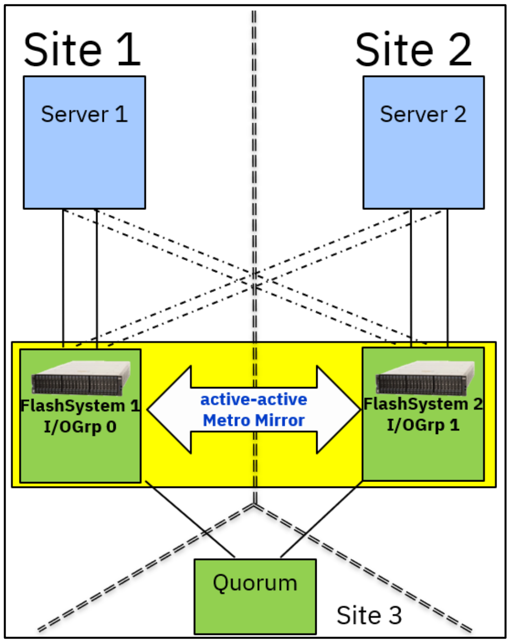

Figure 3: HyperSwap architecture

Business benefits of HyperSwap

HyperSwap enables the SAP HANA application and database server to perform automatic failover of storage access across failure domains with no change in the existing host infrastructure. This ensures continuous availability of services during an unforeseen outage at either site.

The IBM Spectrum Virtualize data efficiency features like deduplication and compression help to reduce the amount of storage that is required to maintain two copies across two failure domains. Using IBM FlashCore modules with data reduction pool can improve the response times of business transactions, with consistent sub-millisecond response times even with data efficiency features in place.

The combination of HyperSwap with IBM FlashSystem can be designed in such a way that the workload is balanced on both failure domains so that none of the FlashSystems are idle during regular operations.

Due to the HyperSwap architecture (active – active two site mirror) writing large blocks of data consumes more time compared to standard single site IBM FlashSystem volumes. However, since the throughput of IBM FlashSystem even with HyperSwap active is many factors above the SAP HANA requirements, there is no operational impact using HyperSwap.

Furthermore, all HANA database read operations and write operations with small block sizes, which are most often used in SAP HANA are not impacted by HyperSwap, as shown in Figure 4. The bars compare the throughput data rates required by SAP HANA KPIs (minimum throughput as required by SAP) with HyperSwap and standard volumes throughput. A larger bar means a higher throughput – better result.

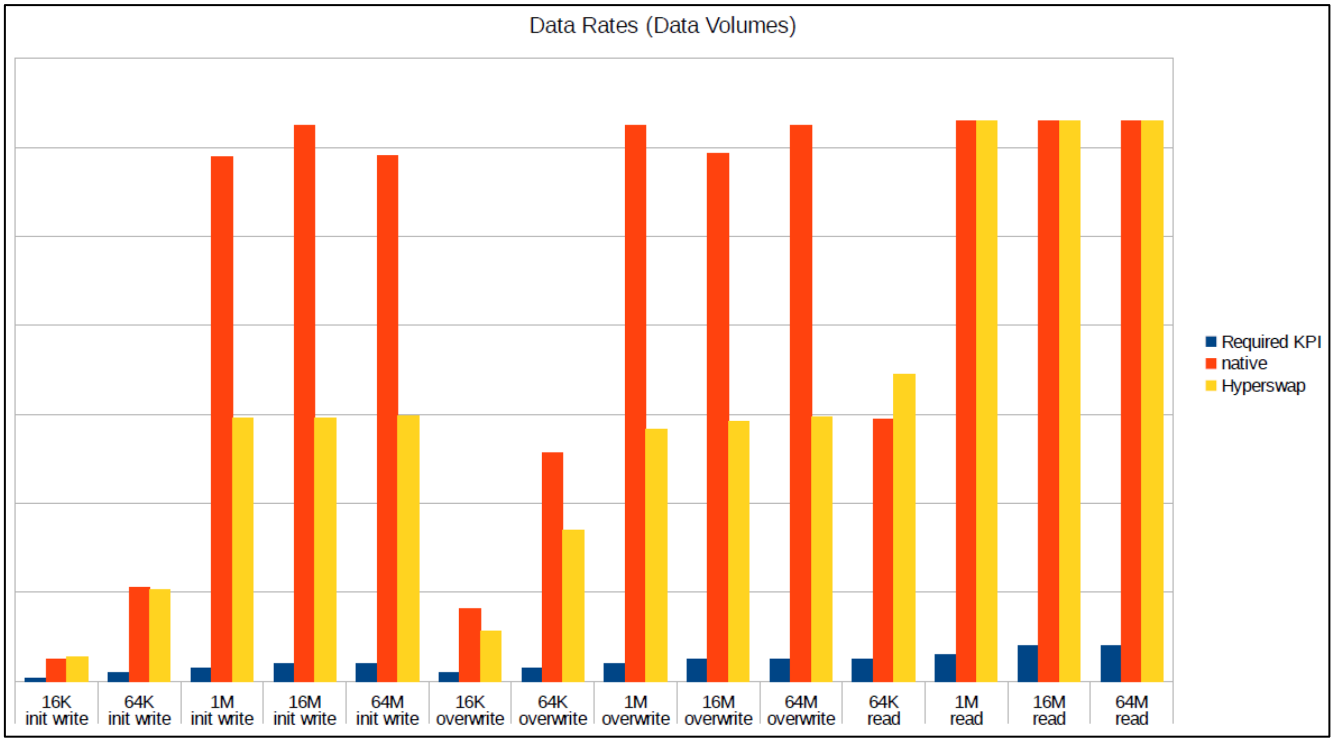

Figure 4: Throughput data rates – KPIs required vs. achieved. Higher is better!

For more information about the IBM FlashSystem family or HyperSwap see the following resources:

- IBM FlashSystem Best Practices and Performance Guidelines: http://www.redbooks.ibm.com/redpieces/abstracts/sg248503.html?Open

- IBM Spectrum Virtualize HyperSwap SAN Implementation and Design Best Practices http://www.redbooks.ibm.com/abstracts/redp5597.html?Open

- IBM HyperSwap: An automated disaster recovery solution https://www.ibm.com/downloads/cas/WJJXG89R

- IBM FlashSystem HyperSwap details in the IBM Knowledge Center: https://www.ibm.com/support/knowledgecenter/STSLR9_8.4.0/com.ibm.fs9200_840.doc/svc_hyperswapovr.html

SAP HANA Setup

For the testing covered by this document the following setup was used, as shown in Figure 5 . The hardware setup consists of:

- 2 x IBM FlashSystem 7200 with:
  - 768 GB System Memory per I/O Group
  - 24 x 4.8 terabytes usable (TBu) / 21.99 terabytes effective (TBe) NVMe IBM FlashCore® Modules (FCM) 2.0
  - 2 x 32Gb FC 4 Port Adapter Pair per node

- IBM Power System H922, 2TB Memory
- 32Gb SAN infrastructure

- LPAR for SAP Netweaver Application
  - 2 cores
  - 256GB memory
  - Operating System: SLES® 12 SP5
  - SAP NetWeaver® 7.5

- LPAR for SAP HANA database
  - 8 cores dedicated
  - 1,5TB memory
  - Operating System: SLES 15 SP1
  - HANA 2.0 SPS05
  - 2 x 32Gb virtual fibre channel ports

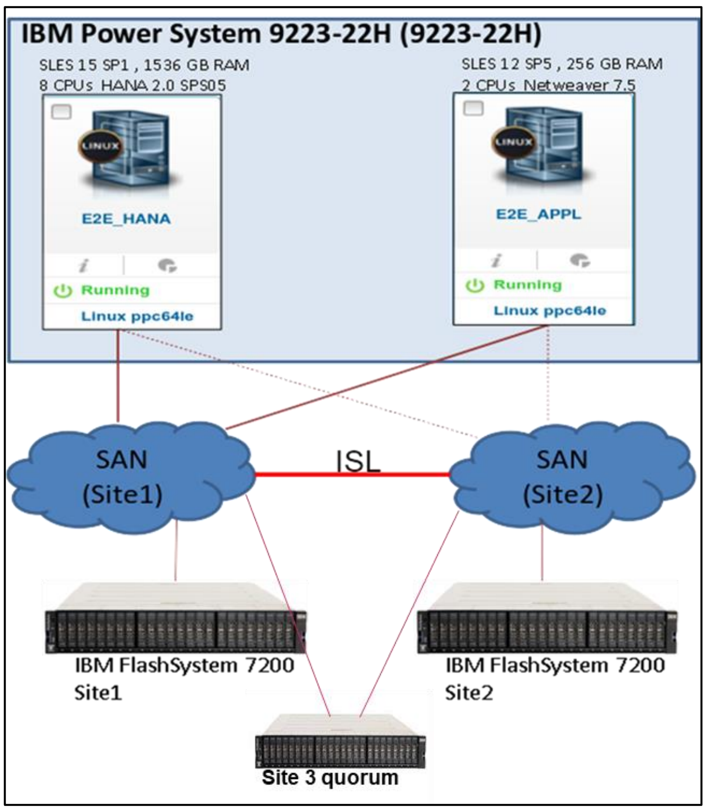

Figure 5: Hardware setup for SAP HANA with IBM FlashSystem HyperSwap testing

The SAP HANA storage configuration followed the recommended guidelines of the IBM System Storage Architecture and Configuration Guide for SAP HANA TDI.

https://www.ibm.com/support/pages/ibm-system-storage-architecture-and-configuration-guide-sap-hana-tdi-v221

Logical configuration of the IBM FlashSystem 7200

On each site we configured a single DRAID 6 pool containing 12 IBM FlashCore Modules on the IBM FlashSystem. The distribution is 9+P+Q, as shown in Figure 6.

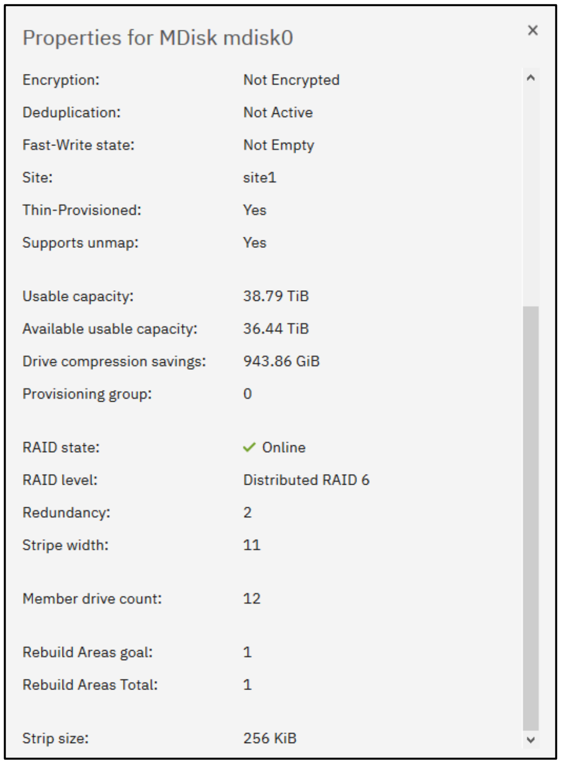

Figure 6: Properties of the configured array (mdisk0)

Test scenarios

The following test scenarios were used for the development of this technical paper:

- SAP HANA KPI tests
- Database startup time
- Application driven business warehouse operations
- Database backup
- Database restores from backup

This set of test scenarios was performed on two different storage configurations. In the first configuration we used *non-HyperSwap* standard volumes and for the second configuration we used *HyperSwap* enabled volumes.

SAP HANA KPI tests

The Storage used for SAP HANA must match a set of minimum performance requirements, so called KPIs (Key Performance Indices). SAP offers an official tool to measure these KPIs, called HCMT.

It allows customers and partners to collect information on the infrastructure intended for SAP HANA deployment. The tool measures whether the planned hardware complies with the requirements defined by SAP. Furthermore, it gauges whether the storage system planned for SAP HANA deployment can achieve satisfactory performance by meeting the minimum Key Performance Indicators (KPIs) requirements as well as satisfactory overall performance given the intended SAP HANA usage. The measurement results are saved into a file, which can be uploaded to the SAP HANA hardware and cloud measurement analysis for further analysis and reporting.

For further information about the SAP HANA HCMT tool, please refer to:

https://help.sap.com/doc/af47cce52aaa4ed4992d42d3cf319d62/2.0/en-US/How_to_Use_the_SAP_HANA_Hardware_and_Cloud_Measurement_Tools_en.pdf

The following test scenarios were used to measure the storage system performance for HANA workload:

Database startup time

SAP HANA is an in-memory database, which requires the complete database loaded into RAM prior starting database operations. In case of a database restart the startup time is determined by two factors:
The database size and the storage read performance.
The startup time is a critical factor, for example in case of an unplanned restart or a site failover.
In this test we measure the duration of the startup from issuing the “HDB start” command to the time when the load has finished.

Application driven business warehouse operations

This test simulates a typical business warehouse (BW) workload: Tables are loaded from an existing database object into a different memory layer. Some kind of analysis processes are then triggered, causing the database to store additional results in new tables. Effectively, this will let the database grow. The expectation is that the SAP HANA savepoints operations will cause additional write traffic on the storage subsystem which we can measure using Storage Insights.

Database backup

The database backup is a second type of write test – a large amount of data is written sequentially to the storage system. Given that the size of an SAP HANA database can be 10TB or more, a high write data-throughput rate is desired when the database is being backed up. To make the scenario more realistic in our tests, we backed-up the database after the load operations had finished as we were interested in the time the backup takes for a database of unchanged size.

Database restores from backup

This is the last sequential read performance test for our storage. Database restores are usually required in case of disaster recovery, or to recover from a logical error. Restore time is always an issue, and it can be drastically shortened by using a fast and high performant storage system.
We always use the same backup database for the restore tests in order to get a fair comparison.

SAP HANA database server storage configuration

In the SAP HANA test environment, the filesystems for storing SAP HANA data volumes and SAP HANA transaction logs are located on LVM logical volumes. For log and for data dedicated LVM volume groups exist. The LVM physical volumes of these volume groups are provided by an IBM FlashSystem 7200.
We measured the IO behavior of the disks from the HANA data and log volume group.

The corresponding data and log VDisks are listed in the IBM FlashSystem Web graphical user interface, as shown in Figure 7.

Name	State	Synchronized	Pool	Protocol Type
e2e_data_1	✓ Online		*Multiple*	SCSI
e2e_data_1 (site1)	✓ Online	Yes	SAP	SCSI
e2e_data_1 (site2)	✓ Online	Yes	SAP_site2	
e2e_data_2	✓ Online		*Multiple*	SCSI
e2e_data_2 (site1)	✓ Online	Yes	SAP	SCSI
e2e_data_2 (site2)	✓ Online	Yes	SAP_site2	
e2e_data_3	✓ Online		*Multiple*	SCSI
e2e_data_3 (site1)	✓ Online	Yes	SAP	SCSI
e2e_data_3 (site2)	✓ Online	Yes	SAP_site2	
e2e_data_4	✓ Online		*Multiple*	SCSI
e2e_data_4 (site1)	✓ Online	Yes	SAP	SCSI
e2e_data_4 (site2)	✓ Online	Yes	SAP_site2	

Figure 7: List of IBM FlashSystem mapped VDisks

Business continuity with Hyperswap

The *HyperSwap* high availability feature available in the IBM Spectrum Virtualize and FlashSystem products enables business continuity during a hardware failure, power outage, connectivity problem, or other disasters, such as fire or flooding. In our setup we simulated a failover by disabling all IBM FlashSystem Fibre Channel ports at Site 1. During the failover a restore job of a 1,5 Tb SAP HANA database was running.

Hyperswap multipath settings

We need to make sure that each host port is connected to each node. IBM FlashSystem volumes should be accessed by a maximum of eight paths. In an optimal Hyperswap configuration we see on the SAP HANA database server, that we have these eight paths per volume, as shown in Figure 8. Only one of the nodes handles IO for a volume under normal conditions. It's the so-called owning node at the primary site. The multipath driver detects this specific topology by using the SCSI ALUA extension and groups the two paths to the preferred node in a separate and prioritized path group.
Note:
IBM recommends specific multipath configurations for IBM Storage Systems. This documentation can be found in the IBM information web page by using the search function. For example, the multipath configuration for IBM FlashSystem 7200 is available here:
https://www.ibm.com/docs/en/flashsystem-7x00/8.4.0?topic=system-settings-linux-hosts

It's recommended to follow the most recent host attachment guidelines!

We will discover how multipathing will change if one site is failing.

```
e2e-hana:~ # multipath -ll mpaths
mpaths (36005076810800273380000000000025d) dm-14 IBM,2145
size=512G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 2:0:7:2 sdcf 69:48  active ready running
| `- 1:0:9:2 sdcx 70:80  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 2:0:6:2 sdbw 68:160 active ready running
  |- 1:0:6:2 sdal 66:80  active ready running
  |- 1:0:7:2 sdau 66:224 active ready running
  |- 1:0:8:2 sdco 69:192 active ready running
  |- 2:0:9:2 sddp 71:112 active ready running
  `- 2:0:8:2 sddg 70:224 active ready running
```

Figure 8: Multipathing in active Hyperswap state

Injecting Fibre Channel connection errors on Site 1

Breaking the Fibre Channel connection to both nodes of Site 1 will cause Hyperswap to initiate a failover.
In the IBM FlashSystem GUI, we observe that the IBM FlashSystem on site1 is offline, as shown in Figure 9.

Figure 9: IBM FlashSystem on Site 1 is off

Now the Hyperswap cluster is using the former target volume at Site 2 for all IO operations.
All required internal steps are transparent to the user, and there is no manual intervention required to get it working. However, some of the changes done by Hyperswap are visible in the IBM FlashSystem GUI, as we will see later.
During the failover we observe the following multipath changes on the SAP HANA database server, as shown in Figure 10.

```
e2e-hana:~ # multipath -ll mpaths
mpaths (3600507681080027338000000000025d) dm-7 IBM,2145
size=512G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 1:0:5:2 sdau 66:224  failed faulty running
| |- 2:0:5:2 sdaz 67:48   failed faulty running
| |- 1:0:6:2 sdy  65:128  active ready running
| `- 2:0:6:2 sdbm 68:0    active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 2:0:4:2 sdj  8:144   failed faulty running
  |- 1:0:4:2 sdi  8:128   failed faulty running
  |- 1:0:7:2 sdah 66:16   active ready running
  `- 2:0:7:2 sdcf 69:48   active ready running
```

Figure 10: Multipathing during failover

Four paths did fail as expected, and because the preferred node is part of site one, both formerly used paths have been replaced by two paths coming from Site 2. As there are still active paths, Host IO to the device will remain and the SAP HANA database restore job can continue.

Recovering from a Site failure

In most cases, recovering from a Site failure is a complex piece of work which requires good planning. Luckily, in our case it's all handled by the Hyperswap cluster which just turn the formerly blocked Fibre Channel ports. Now Site1 is online again and the data which has changed in the meanwhile is copied back from Site 2, as shown in Figure 11. We still can see that the replication direction is going from Site 2 to Site 1.

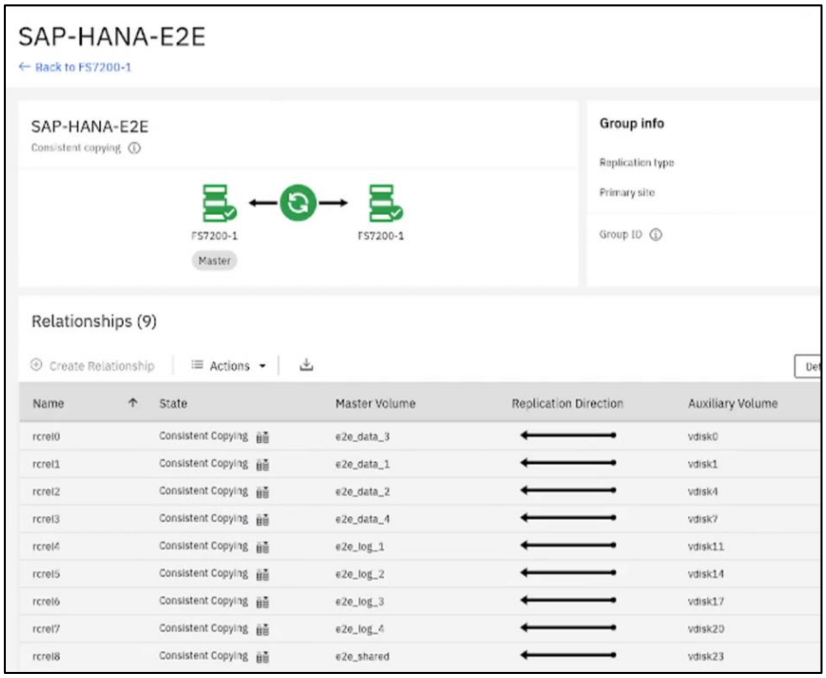

Figure 11: Site 1 is available and changes are copied back

After all changes have been copied back to Site 1, Hyperswap automatically switches the replication direction to the original direction (Site 1 → Site 2) as shown in Figure 12.

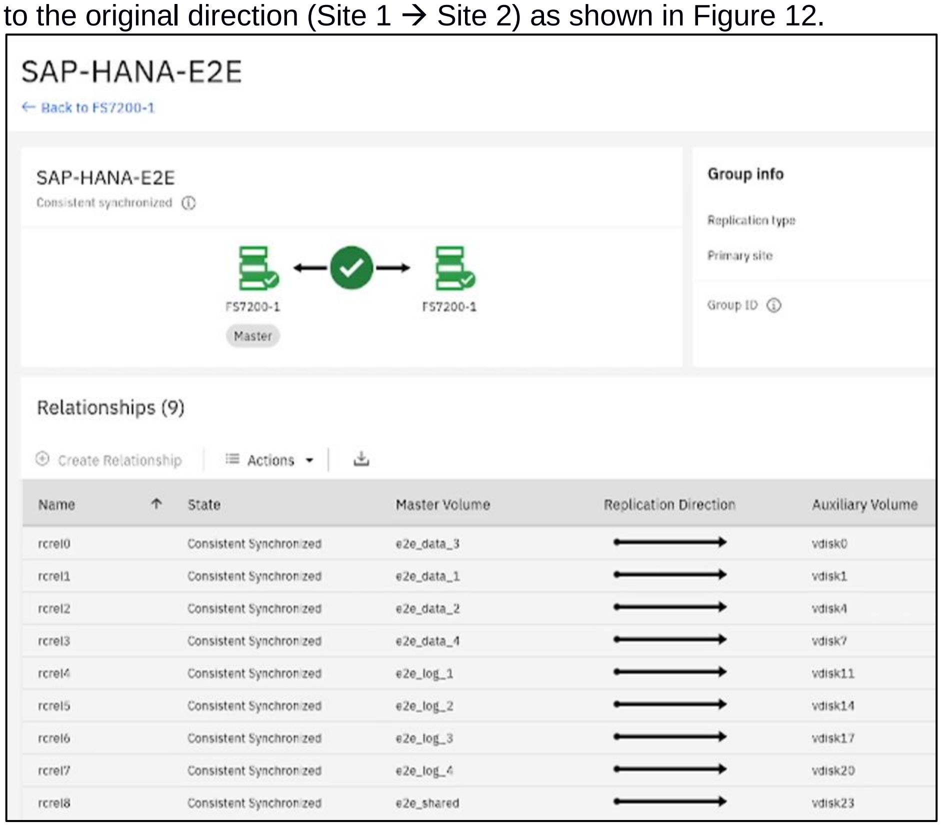

Figure 12: Hyperswap switches back to original copy direction

Results

In this section we describe and analyze the measured results we observed in our test environment.

KPI tests

As described in SAP HANA KPI tests, we used the standard HANA certification measurement tool HCMT to confirm that all required key performance indicators were fulfilled with standard volumes as well as with HyperSwap volumes. Using HyperSwap volumes has no impact on read operations. Due to the active-active Metro Mirror architecture of HyperSwap, write operations are more time consuming. HyperSwap adds some additional latency on write operations with larger block sizes, and the distance between the sites will also add latency to write operations. This is the typical effect of every storage mirroring solution.

HANA database tests

For testing the HANA storage performance two HCMT tests were performed. One with standard volumes and one with HyperSwap volumes. For read operations we saw no difference when using HyperSwap volumes compared to standard volumes. For write operations every mirroring solution will decrease the data throughput, which is also true for HyperSwap. This is especially valid for larger block sizes. However, the SAP HANA KPIs are still overfulfilled when using HyperSwap volumes, as shown in Figure 4 on page 10.

When data is written by the HANA host at either site, both copies are synchronously updated before the write operation is completed. HyperSwap automatically optimizes itself to minimize data that is transmitted between sites, and to minimize host read and write latency.

Due to the HyperSwap architecture, which is called active-active Metro Mirror, both sites are involved in performing write I/Os synchronously. Hence, all FlashSystem nodes on both sites are active, which can be seen in Figure 13.

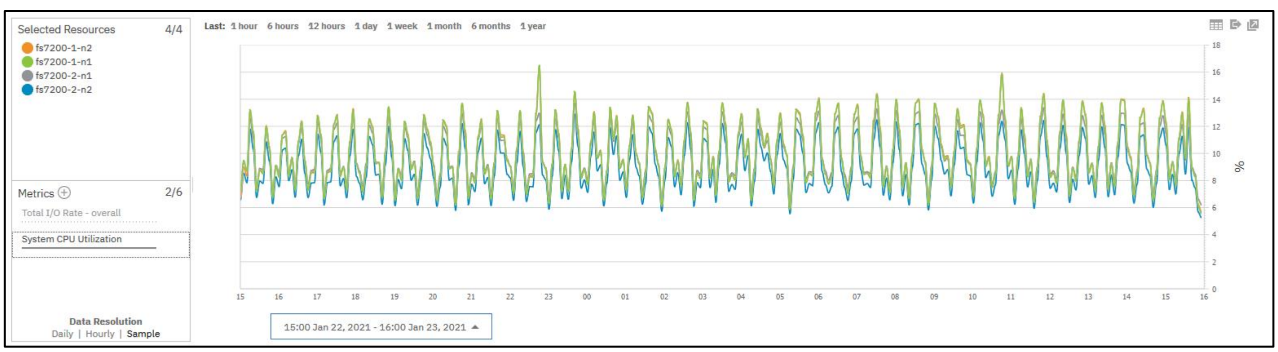

Figure 13: FlashSystem CPU node utilization during business warehouse load operation with HyperSwap

If the same kind of database load runs on non-HyperSwap volumes, only the nodes of the site where the HANA host is attached, shows utilization activity, as shown in Figure 14. Observe the flatlines of the FlashSystem nodes *fs7200-2-n1* and *fs7200-2-n2*.

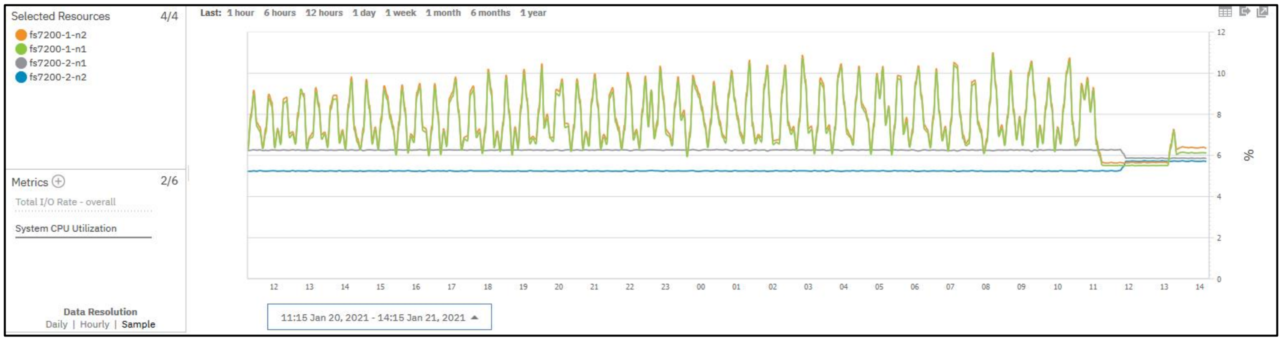

Figure 14: FlashSystem CPU node utilization during business warehouse load operation without HyperSwap.

In addition to the HCMT tests we also tested standard database operations in non-HyperSwap and HyperSwap setup, as shown in Figure 15. For normal business warehouse transactions, the HyperSwap setup has no measurable impact on the database performance. Only disk I/O intensive operations like database backup and restore or initial database startup require slightly more time, due to the HyperSwap architecture.

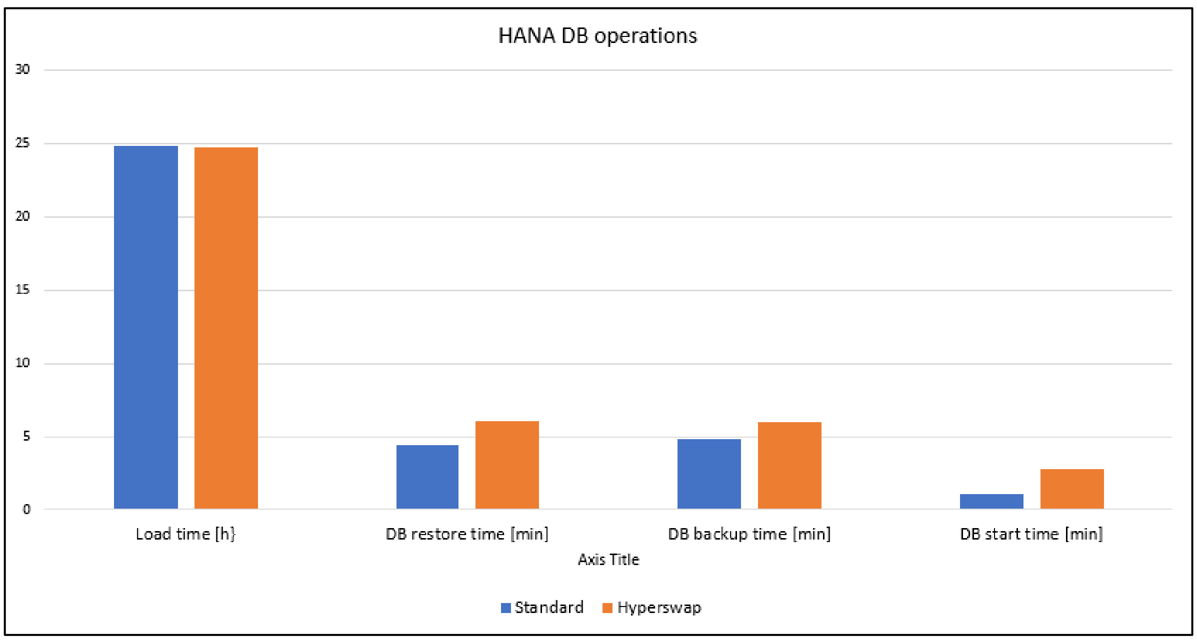

Figure 15: HANA database operations

Why is IBM HyperSwap a preferred solution for high availability?

IBM HyperSwap is built into IBM Spectrum Virtualize software and doesn't require any additional software or hardware components, such as additional multipathing software on the host. It only requires traditional fiber channel (FC or FCIP) connectivity between the sites and native multipathing drivers on the host.

Organizations can also take advantage of IBM's proven storage virtualization features of IBM Spectrum Virtualize to consolidate and unify storage replication requirements at each site.

IBM Spectrum Virtualize also offers advanced data efficiency features such as thin provisioning, compression and deduplication that can be used along with HyperSwap to minimize your total cost of ownership while maintaining two copies across sites.

Summary

Running SAP HANA in TDI environments on IBM FlashSystem provides many benefits like reduced hardware and operational costs, improved availability, and performance. Adding IBM HyperSwap provides continuous availability of the SAP HANA database volumes during planned and unplanned outages of the storage infrastructure. If a FlashSystem at either site goes offline, HyperSwap will automatically fail over storage access to the FlashSystem at the surviving site.

If you are looking for a zero-downtime solution for SAP HANA mission critical workloads, the best solution is to use IBM FlashSystem in combination with IBM HyperSwap.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®
DB2®
Easy Tier®
FlashCopy®
IBM®
IBM FlashCore®
IBM FlashSystem®
IBM Spectrum™
IBM Spectrum Protect™
IBM Spectrum Storage™
IBM Spectrum Virtualize™
POWER®
Power Systems™
POWER9™
PowerLinux™
Redbooks (logo) ®
Storwize®
System Storage®
Tivoli®

The following terms are trademarks of other companies:

Intel, 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.

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

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user, or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details in the section entitled "Cookies, Web Beacons and Other Technologies," and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Back

January 2021