



FlashStack IoT Solution (@SAP COIL)

IoT Design Guide using Virtualized SAP HANA® with Dynamic Tiering on Pure Storage

August 2016

© 2016 Pure Storage, Inc. All rights reserved.

Portions of this documentation are © 2016 SAP SE or an SAP affiliate company. All rights reserved.
Screenshots used with the permission of SAP SE.

Pure Storage, the "P" Logo, Evergreen, FlashStack, and Pure1 are trademarks or registered trademarks of Pure Storage, Inc. in the U.S. and other countries. SAP and SAP Hana are registered trademarks of SAP in the U.S. and other countries. Cisco is a registered trademark of Cisco Systems, Inc. The Pure Storage product described in this documentation is distributed under a license agreement and may be used only in accordance with the terms of the agreement. The license agreement restricts its use, copying, distribution, decompilation, and reverse engineering. No part of this documentation may be reproduced in any form by any means without prior written authorization from Pure Storage, Inc. and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. PURE STORAGE SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

Contributors: Krishna Satyavarapu

Pure Storage, Inc. 650 Castro Street, Mountain View, CA 94041

<http://www.purestorage.com>

Contents

Executive Overview	5
Goals and Objectives.....	5
Audience	7
Design Principles	7
FlashStack Introduction	8
SAP HANA® and Dynamic Tiering–Real time Analytics Platform	9
Deployment of SAP HANA and Dynamic Tiering on FlashStack	11
Cisco UCS Server Configuration.....	13
Cisco UCS Service Profile Configuration	13
Fabric Interconnect Configuration (Cisco UCS 6120XP).....	14
SAP HANA® and Dynamic Tiering	14
Configuration and Optimization of SAP HANA® with Dynamic Tiering	17
Value Propositions of Deploying SAP HANA® with Dynamic Tiering.....	18
Data Reduction.....	18
Summary.....	19
Accelerated Warm Data Readiness.....	19
Summary.....	21
Dashboard Queries.....	22
Summary.....	30
FlashArray//m Sizing for SAP HANA®	31
Throughput Sizing	31
Capacity Sizing	31
Data Volumes Sizing.....	32
Redo Log Volumes Sizing.....	32
Shared Volume Sizing.....	33
Backups.....	33
Total Capacity Size.....	34
Data Reduction of FlashArray//m.....	34
Dynamic Tiering Capacity Sizing.....	34
Conclusions	35
Appendix 1: Pure Storage Components in FlashStack	35
FlashArray //m.....	35
Purity Operating Environment.....	38
Pure1.....	39
Experience Evergreen™ Storage	39
Appendix 2: Cisco Components in FlashStack.....	40
Greater Time-on-Task Efficiency.....	40

Easier Scaling 40

Virtual Blade Chassis..... 40

Reference 41

Executive Overview

The Pure Storage FlashStack™ offering is a flexible, all-flash converged infrastructure (CI) solution from Pure Storage that brings the flash revolution to your data center, faster. It combines the latest in compute, network, storage hardware and virtualization software into a single, integrated architecture that reduces time to deployment, lowers overall IT costs and reduces deployment risk. Highly efficient components reduce the costs associated with power, cooling and data center space. Based on 100% flash-memory storage, CI provides the performance and reliability business-critical IoT applications demand.

The Internet of Things (IoT) is a system of interrelated computing devices, mechanical and digital machines, objects, and people that are provided with unique identifiers which have the ability to transfer data over a network requiring minimal human-to-human or human-to-computer interaction.

SAP® HANA® is an in-memory data platform that is deployable as an on-premise appliance, or on the cloud. It is a revolutionary platform that's best suited for performing real-time analytics, and developing and deploying real-time applications like IoT. Dynamic tiering is an add-on option to the SAP HANA database that enhances the SAP HANA memory store with a flash-backed columnar store. The purpose of SAP HANA dynamic tiering is to offload from SAP HANA memory, the storage and processing of data that has aged (warm data). Building a solution on Pure Storage all-flash enterprise storage will offer significant performance even for warm data and also reducing the overall TCO for the big data application like IoT application. This solves the biggest hurdle enterprises often face when dealing with huge volumes of SAP HANA Data.

In this paper, we will analyze the cost and performance advantage that SAP HANA and dynamic tiering powered by FlashStack offers over traditional database models. We will also run a comparison with maintenance of data on Oracle® (using data reported by SAP) to understand the enhanced performance and reduced total cost of ownership (TCO) presented by dynamic tiering on the FlashStack solution.

Goals and Objectives

Typically, enterprise customers are faced with the two main challenges of maintaining the performance of real-time reporting of business queries and simultaneously reducing the TCO of SAP HANA deployment (especially in big data scenarios like IoT).

The goal of this document is to address these challenges on a highly available and robust Pure Storage infrastructure. The benefits of deploying IoT solution on Pure Storage infrastructure include:

- Up to 50% reduction in deployment time and day-to-day maintenance.
- TCO reduction for SAP HANA—With dynamic tiering model, the need to scale out or scale up SAP HANA storage is eliminated. As volume of data grows, SAP HANA maintenance gets very expensive with licensing and hardware; but with FlashStack and dynamic tiering, storage can be expanded with a marginal increase in cost and with zero compromise on performance.
- Performance consistency—Real-time reporting on SAP HANA® and dynamic tiering deployed on Pure Storage infrastructure is faster for business queries. Given the sub-millisecond latency,

performance is enhanced even when queries do not return data from SAP HANA, but from dynamic tiering storage.

This document describes a Pure Storage reference architecture for deploying an IoT solution using SAP HANA® SPS11 and SAP HANA Dynamic Tiering on a completely virtualized environment using vSphere 5.5. Pure Storage has validated this reference architecture within the SAP COIL lab environment. It includes best practices for the benefit of SAP HANA® Basis administrators, storage administrators, and especially for SAP HANA® modellers and SAP HANA Dynamic Tiering modellers.

Audience

The target audience for this document includes storage architects, data center architects, SAP HANA® Basis administrators, and especially for SAP HANA® modellers and SAP HANA dynamic tiering users who want to implement SAP HANA® big data scenarios like IoT systems. A working knowledge of server, storage, networks and data center design is assumed but is not a prerequisite to read this document.

Design Principles

The guiding principles for implementing this reference architecture are:

- **Repeatable:** Create a scalable building block that can be easily replicated at any customer site. Publish the version of various firmware under test and weed out any issues in the lab before customers deploy this solution.
- **Available:** Create a design that is resilient and not prone to failure of a single component. For example, we include best practices to enforce multiple paths to storage, multiple NICs for connectivity, and high availability (HA) clustering.
- **Efficient:** Build a solution that leverages efficient benefits of an all-flash architecture.
- **Simple:** Simplify deployment and ongoing maintenance tasks via automation.
- **Scalable:** Create a design that can start small, but can easily grow to meet the needs of a growing enterprise.

FlashStack Introduction

The FlashStack™ solution is a flexible, all-flash converged infrastructure solution that brings the flash revolution to SAP systems. It combines best-in-class compute and networking components from Cisco® with Pure Storage all-flash storage into an integrated architecture that speeds time to deployment, lowers costs and reduces deployment risks. Based on 100% all-flash storage, The FlashStack offering provides the performance and reliability your business applications demand.




Figure 1: Pure Storage FlashStack

FlashStack CI is available from accredited FlashStack Partners who help provide an excellent converged infrastructure ownership experience. FlashStack Partners have the knowledge and experience necessary to help streamline the sizing, procurement, and delivery of your entire system.

Key Benefits of the FlashStack solution are:

- **Consistent Performance and Scalability**

- Consistent sub-millisecond latency with 100% flash storage.
- Consolidate 100s of enterprise-class applications in a single rack.
- Scale easily, without disruption.
- Repeatable growth through multiple FlashStack CI deployments.

- **Operational Simplicity**

- Fully tested, validated, and documented for rapid deployment.
- Reduced management complexity
- Auto-aligned 512b architecture eliminates storage alignment headaches
- No storage tuning or tiers necessary

- **Lowest TCO**

- Dramatic savings in power, cooling and space with 100% Flash.
- Industry leading data reduction.
- Controller upgrades included with three year maintenance renewal under the Pure Storage Evergreen™ Storage model.

- **Enterprise Grade Resiliency**

- Highly available architecture and redundant components
- Non-disruptive operations
- Upgrade and expand without downtime or performance loss
- Native data protection: snapshots and replication

To know more about the various components in FlashStack, see [Appendix 1: Pure Storage Components in FlashStack](#), [Appendix 2: Cisco Components in FlashStack](#).

SAP HANA® and Dynamic Tiering–Real time Analytics Platform

SAP HANA®, which combines row-based and column-based database technology, is an in-memory data platform that is deployable as an on-premise appliance, or on the cloud. It is a revolutionary platform that is best suited for performing real-time analytics, and developing and deploying real-time applications. Being a combination of hardware and software, it is made to process massive real time data using in-memory computing (an in-memory database means all the data is stored in memory or the RAM). Although data resides in main-memory (RAM), a robust enterprise storage system will enable operations such as persistence and delta merges to operate at maximum efficiency. The use of multi-core CPUs, and multiple CPUs per board, and multiple boards per server appliance accelerates the speed advantages offered by this RAM storage. Complex calculations on data are not carried out in the application layer, but are moved to the database.

SAP HANA® is equipped with multi-engine query processing environment which supports relational, graphical, and text data within same system. It provides features that support significant processing speed, handle huge data sizes and text mining capabilities.

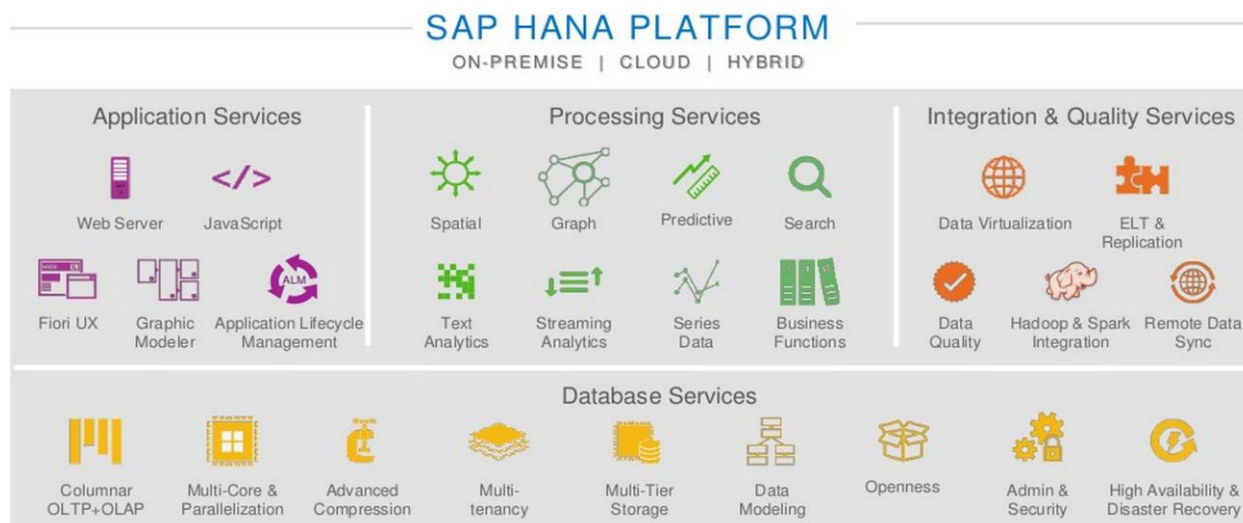


Figure 2: SAP HANA® Platform

Dynamic tiering is an add-on option to SAP HANA database, and is a native big data solution. Dynamic tiering allows for warm data to be moved to extended storage which results in reduced size and better performance of in-memory database.

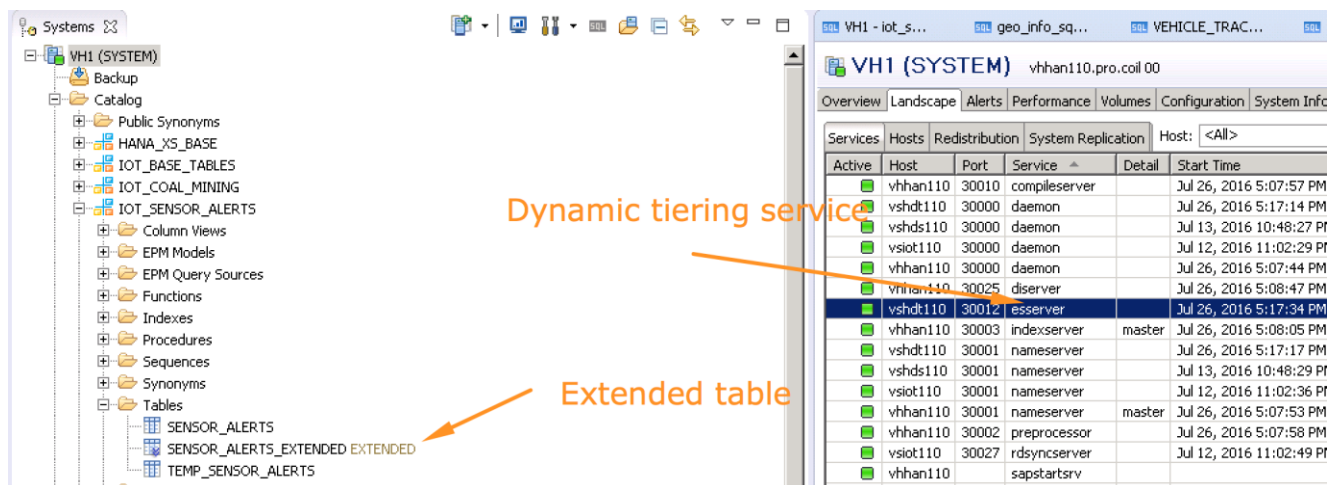


Figure 3: SAP HANA Dynamic Tiering Tables

Dynamic tiering adds the esserver service to your SAP HANA system. This service creates an extended store and extended tables. While extended tables behave like all other SAP HANA tables, their data resides in the storage, in this case, all-flash-based FlashStack. This offloads storage and processing of data that has aged (warm data) from the core HANA memory. SAP Data Lifecycle Manager (DLM), a SAP Data Warehousing Foundation tool, is a HANA application that provides GUI to define data aging rules which can be used to determine the data to be moved to dynamic tiering. It creates stored procedures to perform mass data movement between SAP HANA and dynamic tiering stores and also let you schedule those procedures for automation execution.

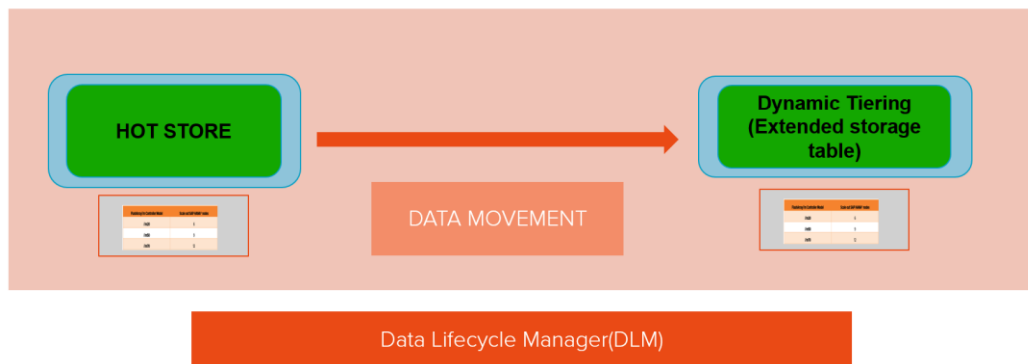


Figure 4: SAP HANA to DT Data Transfer

HANA dynamic tiering solution built on Pure Storage all-flash enterprise storage boosts performance of warm data as well reduces the overall TCO for big data applications like IoT. In big data application environments, as the volume of data grows, storing all the data in SAP HANA could turn out to be expensive. In order to reduce TCO, data management must be done using a dynamic tiering model. Data reporting can be done on dynamic tiering data as well, as the extended tables behave like all other SAP HANA tables, except that their storage location is different. SAP Lumira is used for doing long term reporting on SAP HANA and dynamic tiering. SAP Lumira is a self-service solution that allows analysts and decision makers to access, transform, and visualize data. The SAP Lumira desktop experience is used to prepare data from multiple sources, visualize it, and then compose stories from those visualizations that can be shared with other decision makers. Using the SAP Lumira server and cloud

platforms (which provide browser- and mobile-based experiences), you can further analyze data and collaborate with colleagues on datasets, stories, and other business intelligence artifacts.




Figure 5: SAP Data Reporting

Deployment of SAP HANA and Dynamic Tiering on FlashStack

FlashStack for SAP HANA consists of a combined stack of hardware (storage, network, and compute) and software (Cisco UCS Manager®, SAP HANA® SPS11 and SAP HANA Dynamic Tiering, SAP Design Studio, SAP Smart Data Streaming, SAP Lumira server, and Pure Storage GUI)

- **Network:** Cisco UCS Fabric® Interconnect 6120XP for external and internal connectivity of IP and FC network.
- **Storage:** Pure Storage FlashArray//m20 with Fibre Channel connectivity
- **Compute:** Cisco® UCS B250 M2 Server




Figure 6: FlashStack Connectivity

Figure 6 shows a detailed topology of the reference architecture configuration. A major goal of the architecture is to build a highly redundant and resilient infrastructure. We used powerful servers with dual Fibre Channel ports connected redundantly to two Fabric Interconnects switches that were connected to redundant FC target ports on the FlashArray//m. The servers also have redundant network connectivity. This is configured on switch mode. The SAP HANA and dynamic tiering is completely virtualized using vSphere 5.5.

Cisco UCS Server Configuration

A pair of Cisco UCS Fabric Interconnects 6120XP, and one Cisco UCS B-series B250-M2 blade server were deployed for hosting SAP HANA® SPS11, SAP HANA Dynamic Tiering, SAP Design Studio, SAP Smart Data Streaming and SAP Lumira server. The UCS manager, UCS Fabric Interconnects and the components in the chassis were upgraded to 2.2.3f firmware level.

The server has Cisco UCS M81KR/M71KR-Q cards and they were connected through four ports from each Cisco Fabric extender of the Cisco UCS chassis to the Cisco Fabric Interconnect to access the Pure Storage FlashArray//m LUNs. The server configuration is described in Table 1.

Table 1: Cisco UCS Server Configuration

Component	Description
Processor	2X Intel(R) Xeon(R) 5680 @ 3.333GHz, 12 Cores
Memory	256 GB @ 1333 MHz, regular voltage
HBA	4 X 16G ports on Cisco UCS M81KR/M71KR-Q
NIC	2 X 10G ports on Cisco UCS M81KR/M71KR-Q
Application	SAP HANA Dynamic Tiering , SAP Design Studio, SAP Smart Data Streaming, SAP Lumira server
Database/Platform	SAP HANA® SPS11
Virtualization	vSphere 5.5

Cisco UCS Service Profile Configuration

In order to facilitate rapid deployment of UCS servers, we created a service profile template with the following characteristics.

- Boot configuration from SAN policy (in order for the server to boot up from a Pure Storage boot LUN)
- Ethernet and FC adapter policy set to VMware
- B250-M2 blade servers set to BIOS defaults
- Two vHBA FC adapters and four vNIC ethernet adapters configured on the Cisco VIC cards to avoid any single point of failure
- Deployment of service profiles from the template with blade servers associated
- All other settings at default

Fabric Interconnect Configuration (Cisco UCS 6120XP)

The Cisco UCS 6120XP 20-Port Fabric Interconnect is a core part of the Cisco Unified Computing System. Typically deployed in redundant pairs, Fabric Interconnects provide uniform access to both networks and storage. Here we are configuring it for the Switch Mode.



Figure 7: Cisco UCS 6120XP

SAP HANA® and Dynamic Tiering

The setup simulated for this test scenario comprised the below configuration:

- A single node, completely virtualized environment using vSphere 5.5.
- One Cisco UCS B250 M2 blade used with SUSE 11 SP3 OS with SAP HANA Dynamic Tiering, SAP Design Studio, SAP Smart Data Streaming, SAP Lumira server, and FlashArray//m20.
- Four 16GB Fibre channel ports from the array attached to the SAN Fabric; each node with two fibre channel ports for storage and four 10 GBPS ports for network.

The following table outlines the SAP software matrix:

Table 2: Software Matrix used for SAP HANA® system

S/no	Component	Software
1.	Database/Platform	SAP HANA® SPS11, Dynamic tiering, Smart data streaming
2.	Operating System	SUSE® Linux® Enterprise 11 Service Pack 3
3.	Reporting Application	SAP Design Studio 16.2.1, SAP Lumira server 1.30
4.	File System	XFS

All the servers boot from SAN, with dedicated drive for their OS related partitions. All other LUNs are shared across the nodes. The /hana/shared partition is used for storing SAP binaries and is mounted/shared in parallel simultaneously across SAP HANA, dynamic tiering, and smart data streaming.

For the purpose of this design guide, we focus on IoT where real-time, high velocity data is received from sensors attached to a Raspberry Pi and also through sensor data simulator, as shown below. This high velocity is first fed to smart data streaming service in SAP HANA and is then sent to SAP HANA using ODBC connector. The data is later transferred to dynamic tiering from SAP HANA.

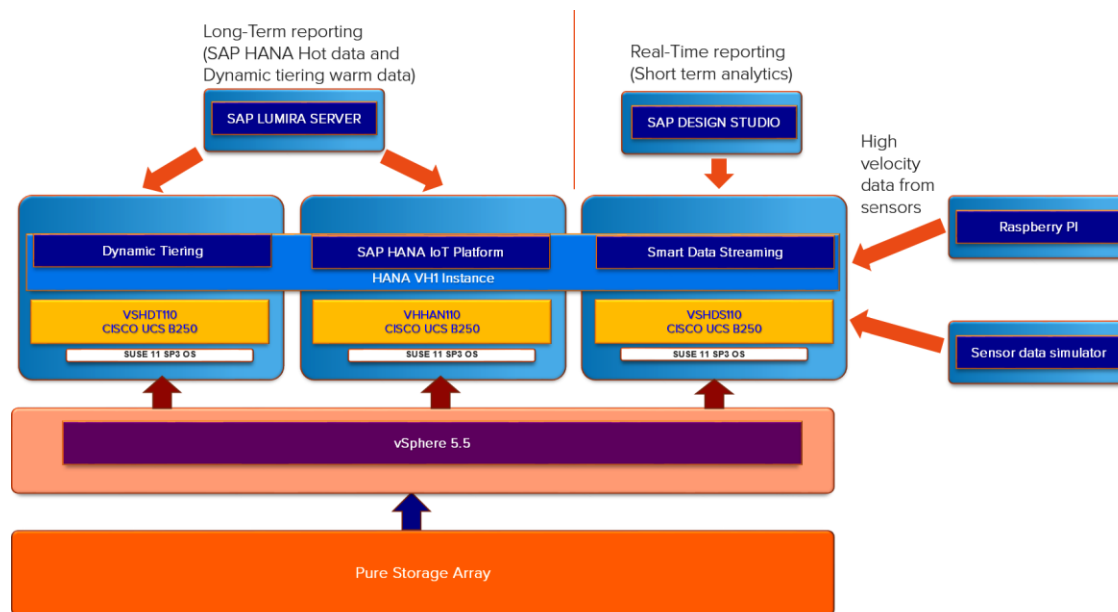


Figure 8: SAP HANA and Dynamic Tiering Setup

We used SAP Design studio for reporting this high velocity data, in real time, with the smart data streaming window as data source. SAP BusinessObjects™ Design Studio is an alternative to BEx Web and BEx Web Application Designer. SAP BusinessObjects Design Studio allows for intuitive design of centrally governable analytic content ranging from guided analytics to sophisticated OLAP applications and aggregated dashboards. Some of its features include out-of-the box iPad support, a state-of-the-art HTML5 UI, seamless application theming, a WYSIWYG Eclipse-based designer, full and native support of BW BEx queries, direct connectivity to HANA, and an advanced scripting engine.




Figure 9: Logical architecture of the Pure Storage HANA® setup

The HANA dynamic tiering (DT) server must be installed in its own server, which is called an extended storage host. The HANA host and extended storage host or dynamic tiering are required to share the binaries stored in `/hana/shared` folder leveraging file systems like NFS (Network File System). Dynamic tiering will use the same directories for data and log backups as SAP HANA. The base paths to the backup directories should therefore be shared.

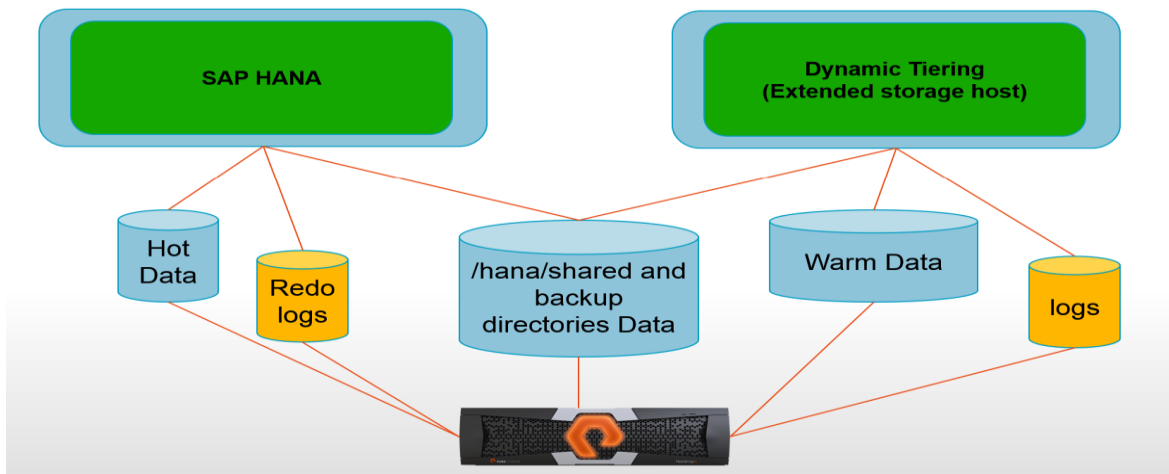


Figure 10: SAP HANA and Dynamic Tiering File System

Configuration and Optimization of SAP HANA® with Dynamic Tiering

In this section, we will cover the configuration and optimization for dynamic tiering. The settings are in the configuration file called esserver.ini. Here are the some of the important parameters:

Table 3: Configuration and Optimization of SAP HANA® with Dynamic Tiering

Sl.No.	Name of the parameter	Description	Recommended settings
1.	catalog_cache	Amount of memory initially reserved for caching the dynamic tiering catalog.	The catalog cache should be sized at 2 to 8 times the size of the catalog file. The catalog file contains metadata for extended tables. The file has the name "<HANA_System_ID>ESDB.db", and can be found in a subdirectory rooted at /hana/data_es/<HANA_System_ID>.
2.	delta_memory_mb	Amount of memory available to store delta-enabled extended tables.	Size this cache only large enough to keep data in memory that is being loaded and changed prior to being merged with the main store data on disk. This amount will vary depending on the number of tables, load frequency, and merge frequency with the main store on disk. A good starting point for most systems is 250 - 500MB for each delta-enabled extended table. This should allow for enough of a safety margin in case load frequencies spike on all objects simultaneously.
3.	load_memory_mb	Maximum amount of memory extended storage can request from the operating system for temporary use.	Calculate 85% of the total memory on the DT server machine, subtract the amount of memory you will need for RLV memory (delta_memory_mb), and then divide the remaining memory in thirds and give a third to each cache: main_cache_mb, temporary_cache_mb, and load_memory_mb
4.	main_cache_mb	Amount of memory to be used for caching dynamic tiering database objects.	Calculate 85% of the total memory on the DT server, subtract the amount of memory you will need for RLV memory (delta_memory_mb), and then divide the remaining memory in thirds and give a third to each cache: main_cache_mb, temporary_cache_mb, and load_memory_mb

Value Propositions of Deploying SAP HANA® with Dynamic Tiering

In this section, we will know more about the value that SAP HANA and Dynamic Tiering on FlashStack offers. The major points of focus are:

- Data reduction
- Accelerated warm data readiness
- Dashboard queries
- Cost Analysis

Data Reduction

Dynamic tiering is a Sybase IQ system, a columnar disk-based relational database. Dynamic tiering has similar characteristics of data compression to SAP HANA. Sybase IQ compresses in bitmaps and SAP HANA compresses with dictionary encoding.

To test data reduction on dynamic tiering, we used 200GB of TPC-H data. We initially loaded all the tables as column stores and then manually moved the data to dynamic tiering using the following SQL command.

```
INSERT INTO <Extended Table> SELECT * FROM <HANA_table>
```

Initially with all the 200GB column store (100% column store), data reduction on the dynamic tiering data volume was between **2.2 and 2.6**.

Table 4: Data Reduction Comparison

Data Loaded	Minimum data reduction	Maximum Data Reduction
SAP HANA	1.9	2.3
Dynamic tiering	2.2	2.6




Figure 11: Data reduction Analysis Graph

Summary

The above results prove that FlashArray//m enhances data reduction of SAP HANA and Dynamic tiering further. We note that the data reduction numbers of dynamic tiering were better than SAP HANA for the same set of data.

Accelerated Warm Data Readiness

Time taken for warm data to move to SAP HANA is an important factor because the set of tables will be on downtime during the transfer and cannot be queried.

In this case we used customized SQL to move the data between SAP HANA and dynamic tiering manually as access to the DLM tool was unavailable. We then reported data by creating an association between the two data stores.

In order to test this, we created a table with around 300 bytes of structure width in SAP HANA and populated it from 50 million records to 500 million records. We moved the data to dynamic tiering using the following SQL command:

```
INSERT INTO <Extended Table> SELECT * FROM <HANA_table>
```

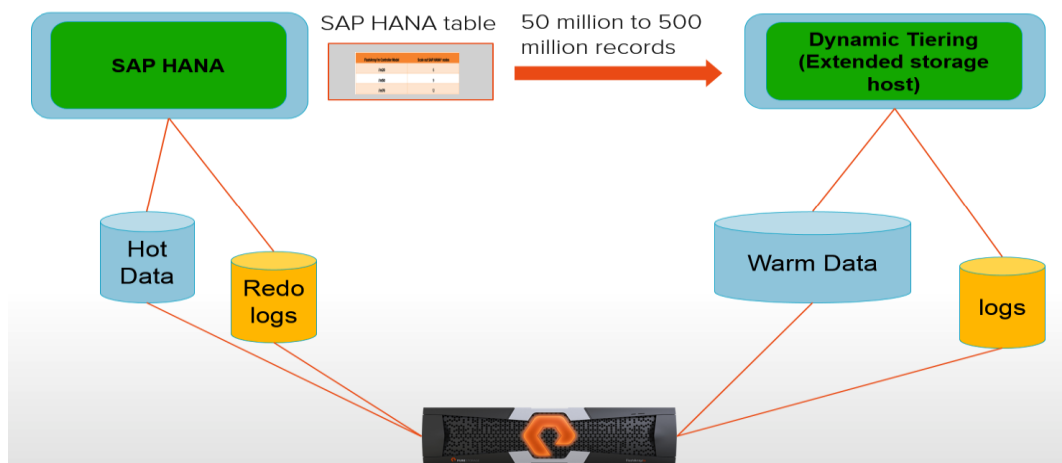


Figure 12: Data Transfer from SAP HANA using SQL Command

As you can see below, the data movement on FlashArray//m is extremely fast due to its superior write performance. With better CPU/memory performance, better write performance can be observed.

Table 5: Data Transfer from SAP HANA - Performance Analysis

Number of Records	Time taken to move data to DT	Peak bandwidth achieved	Avg. bandwidth achieved
50 million	1 min 3 seconds	63 MB/sec	40 MB/sec
250 million	4 mins 17 seconds	278 MB/sec	165 MB/sec
500 million	7 mins 23 seconds	385 MB/sec	229 MB/sec

The following graph depicts is the Pure Storage GUI view of the movement of 250 million records from SAP HANA to dynamic tiering. This shows our write performance and the rate at which we can push the warm data to dynamic tiering and make it available for reporting, or the rate at which we can free up SAP HANA memory. As you can see, we were able to achieve write bandwidth of **165 MB/second** with **278 MB/second peak bandwidth** and entire data move to dynamic tiering took little more than 4 minutes.




Figure 13: Transfer of 250 Million Records from SAP HANA

Below is the analysis of peak/average bandwidth achieved with the data movement to dynamic tiering from SAP HANA.




Figure 14: Peak/Average Bandwidth during Data Transfer

Summary

Due to the superior write performance of the FlashStack solution, data movement to dynamic tiering is accelerated, which in turn reduces the overall downtime of accessing the tables.

Cost and Performance Analysis

As SAP HANA data volume grows, it needs to be scaled-out or scaled-up. Both options could turn out to be expensive, given the combined cost of SAP HANA and the associated hardware. However, with the robust read performance of Pure Storage FlashArray//m, the performance of queries on extended tables will be comparable to those on SAP HANA. This also offers a TCO advantage, since cost of licensing and hardware cost of dynamic tiering is very less compared to that of SAP HANA.

Below is the detailed analysis of how we tested this scenario and cost per performance analysis.

Dashboard Queries

We used the IoT table in SAP HANA called `SENSOR_ALERTS` (300 bytes) in which large volumes of data gets pumped from smart data streaming. We moved this data to dynamic tiering where we had created an extended table called `SENSOR_ALERTS_EXTENDED` (with the same structure as `SENSOR_ALERTS`).

On these tables `SENSOR_ALERTS` and `SENSOR_ALERTS_EXTENDED`, we created a calculation view (shown below) and performed a union of the SAP HANA table and dynamic tiering table. Creation of calculation view is possible as dynamic tiering tables behave exactly like SAP HANA tables. Calculation view aids the comparison of performance of business queries running on SAP HANA and dynamic tiering tables directly.

Initially, we populated the SAP HANA tables with > 500 million records. We then scaled down the number of records to zero by pushing the data to dynamic tiering table, and scaled dynamic tiering up from 0 to 500 million records. We ran the three queries as shown in the SAP Lumira dashboard below, with all of them doing 500 million records scan.

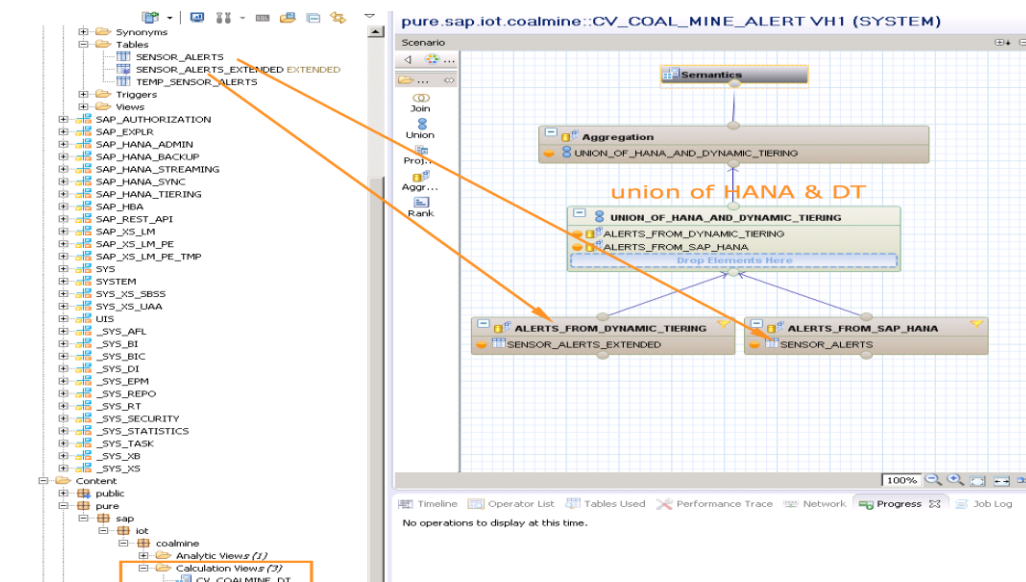


Figure 15: Calculation View and Union of SAP HANA and DT Tables

Below is the breakdown of SAP Lumira dashboard queries hitting SAP HANA and dynamic tiering calculation view. This dashboard is running on the calculation view → `pure.sap.iot.coalmine/CV_COAL_MINE_ALERT`, which is based on the union of tables based in SAP HANA \and dynamic tiering `SENSOR_ALERTS` and `SENSOR_ALERTS_EXTENDED`.




Figure 16: SAP Lumira Dashboard Queries on Calculation View

This is the breakdown of the queries. As shown below, all the queries perform a 500 million record scan.

Table 6: IoT Queries

Query ID	Query (From SAP HANA® SQL Plan Cache)	Calculation View	Total records
Q1(Detailed analysis of alerts by city)	SELECT "CITY", "ALERT_TYPE", SUM("ALERT_COUNT") AS "ALERT_COUNT_SUM" FROM "_SYS_BIC"."pure.sap.iot.coalmine/CV_COAL_MINE_ALERT" GROUP BY "CITY", "ALERT_TYPE"ORDER BY "CITY" ASC, "ALERT_TYPE" ASC	pure.sap.iot.coalmine/CV_COAL_MINE_ALERT	500 million records
Q2(Aggregated alerts per state)	SELECT "STATE", SUM("ALERT_COUNT") AS "ALERT_COUNT_SUM" FROM "_SYS_BIC"."pure.sap.iot.coalmine/CV_COAL_MINE_ALERT" GROUP BY "STATE" ORDER BY "STATE" ASC	pure.sap.iot.coalmine/CV_COAL_MINE_ALERT	500 million records
Q3(Alert count vs Alert Type)	SELECT "ALERT_TYPE", SUM("ALERT_COUNT") AS "ALERT_COUNT_SUM" FROM "_SYS_BIC"."pure.sap.iot.coalmine/CV_COAL_MINE_ALERT" GROUP BY "ALERT_TYPE" ORDER BY "ALERT_TYPE" ASC	pure.sap.iot.coalmine/CV_COAL_MINE_ALERT	500 million records

Case 1—100% Data in SAP HANA / 0% data in DT (Baseline test case): These queries were run when all the data was present in SAP HANA and no data present in DT.

Query Name	Query Description	Number of Records in SAP HANA	Number of Records in DT	Total Runtime
Q1	Detailed analysis of alerts by city	~500 million records	0	0.879 second
Q2	Aggregated alerts per state	~500 million records	0	0.643 second
Q3	Alert count vs Alert Type	~500 million records	0	0.547 second

Case 2—75% Data in SAP HANA / 25% data in DT: These queries were run when 25% of data was moved to DT, which is around 125 million records.

Query Name	Query Description	Number of Records in SAP HANA	Number of Records in DT	Total Runtime
Q1	Detailed analysis of alerts by city	~ 375 million records	~125 million records	1.604 seconds
Q2	Aggregated alerts per state	~ 375 million records	~125 million records	.1318 second
Q3	Alert count vs Alert Type	~ 375 million records	~125 million records	1.172 second

Case 3—50% Data in SAP HANA / 50% data in DT: These queries were run when 50% of data was moved to DT, which is around 250 million records.

Query Name	Query Description	Number of Records in SAP HANA	Number of Records in DT	Total Runtime
Q1	Detailed analysis of alerts by city	~250 million records	~250 million records	2.197 seconds
Q2	Aggregated alerts per state	~250 million records	~250 million records	1.730 seconds
Q3	Alert count vs Alert Type	~250 million records	~250 million records	1.523 seconds

Case 4—25% Data in SAP HANA / 75% data in DT: These queries were run when 75% of data was moved to DT which is around 375 million records. Here are the results of the same.

Query Name	Query Description	Number of Records in SAP HANA	Number of Records in DT	Total Runtime
Q1	Detailed analysis of alerts by city	~125 million records	~ 375 million records	2,671 seconds
Q2	Aggregated alerts per state	~125 million records	~ 375 million records	2.368 seconds
Q3	Alert count vs Alert Type	~125 million records	~ 375 million records	2.189 seconds

Below is the graphical analysis showing the comparison of dashboard queries runtime performance when the data is in SAP HANA and when data is gradually moved to DT. We notice that the query runtimes do not vary sharply when compared to baseline runtime when data increases or decreases in DT. This shows the superior read performance of the FlashArray//m.

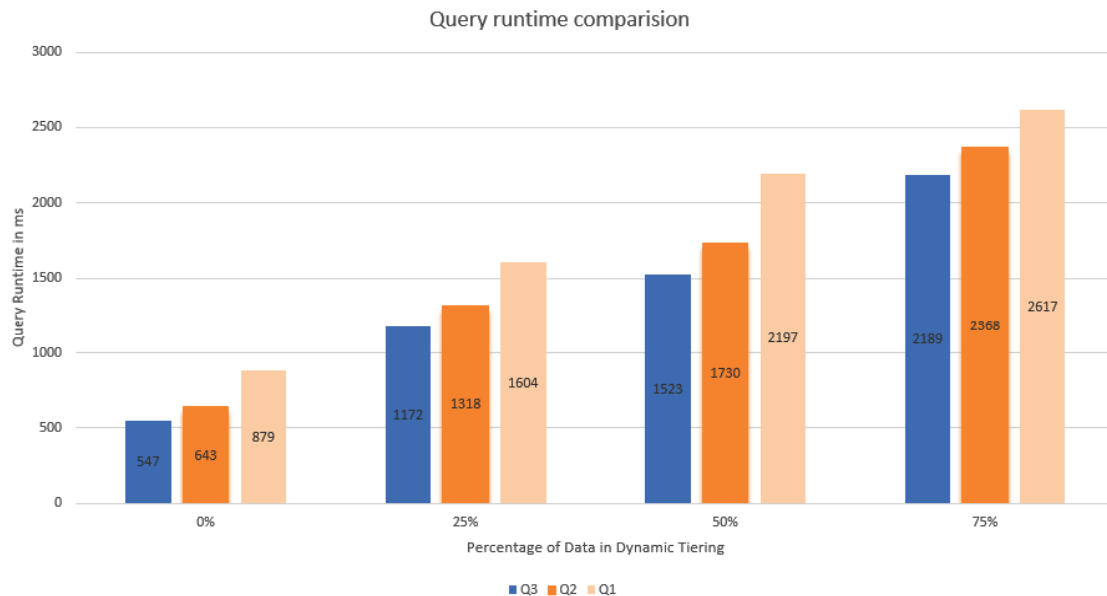


Figure 17: Query Runtime Comparison Graph

Cost and Performance Analysis

Let us compare the cost associated with maintenance of data, both in SAP HANA and in dynamic tiering (both running on the FlashStack™ solution) in the backdrop of increase of volume of data over time. In this use case scenario, we start with 1 TB of data in SAP HANA with a data increase of 250 GB over a period of time.

First, let us analyze the cost of deployment of SAP HANA with and without dynamic tiering. We will then compare the performance of queries on Oracle database and on SAP HANA with dynamic tiering. The performance analysis on Oracle database is based on SAP community public site mentioned below:

<https://scn.sap.com/people/sheetal.jain/blog/2015/05/21/hana-performance-test-benchmarks>

Case 1—Data load on SAP HANA DT increments by 25%

In this test scenario, we will compare the performance of query runtime on SAP HANA DT with results reported by SAP compared against Oracle. We tested the cost and performance advantage by reducing the data on SAP HANA by 25% and increasing the DT load correspondingly. We will then take a combined view of the query runtime and cost advantage of the data on SAP HANA DT with that (as reported by SAP) of data on Oracle database, based on findings from the SAP HANA public site¹. The license cost of SAP HANA is \$150,000 per 64 GB of data (per the references cited ²).

SAP HANA with DT cost and performance advantage as compared to Oracle

From the following table, we observe that even when just 25% of data was moved to DT, despite a marginal drop in performance of query runtime, the cost still proved to be advantageous by savings of around \$750,000. Also, as based on the SAP reported data, when compared to Oracle, the performance of query runtime of data on DT was still 9 times better.

Table 7: Cost and Performance Comparison (HANA Vs Oracle), using SAP reported data for Oracle (see References)

Case	1	2	3	4
Initial SAP HANA Size	1TB			
HANA (data volume %)	100	75	50	25
Dynamic tiering (Data volume %)	0	25	50	75
~Query time (seconds)	0.69	1.36	1.82	2.4
Variation (seconds)	0.00	1.97	2.64	3.48
Performance difference (HANA vs Oracle)	13.90	11.93	11.26	10.42
HANA License Cost [Millions] \$150k per 64GB	\$ 2.484	\$ 1.863,28	\$ 1.242,19	\$ 621,09
Maintenance [Millions] 22% per year	\$ 0.546	\$ 0.409	\$ 0.273	\$ 0.136
Total cost (Millions)	\$ 3.030	\$ 2.273	\$ 1.515	\$ 0.757
Savings in Millions	\$ -	\$ 0.757	\$ 1.515	\$ 2.273

¹ <https://scn.sap.com/people/sheetal.jain/blog/2015/05/21/hana-performance-test-benchmarks>

² <http://www.asugnews.com/article/sap-hana-pricing-amid-controversy-sap-offers-licensing-clarity>
<http://www.kernelsoftware.com/products/catalog/sap.html>

The graph below shows the performance difference between Oracle and SAP HANA with DT for queries. This shows that queries run 6 – 13.9 times faster on SAP HANA with DT compared to Oracle, using the SAP reported data for Oracle. Besides this, there is a significant savings of spending per year in terms of SAP HANA license cost and maintenance.




Figure 18: Performance Comparison between SAP HANA with DT and Oracle

SAP HANA with DT TCO advantage as compared to SAP HANA Scale Up Strategy

The following graph illustrates the economy that SAP HANA and DT model offers when compared with scale up model of SAP HANA, based on 1 TB size. This accounts for savings of cost in terms of SAP HANA licensing and maintenance as more data is moved to dynamic tiering by reducing the SAP HANA footprint.



Figure 19: Cost Comparison between SAP HANA with DT and SAP HANA Scale Up

Case 2—Volume of data on SAP HANA grows by 50%

In this test scenario, data grows by 50% in SAP HANA from 1 TB to 1.5 TB. The following table highlights the cost advantage by comparing the cost of keeping 100% data in SAP HANA and the cost (and associated saving) of moving it to dynamic tiering. You will notice that with even 25% move to DT from SAP HANA, the saving is as much as one million dollars per year by way of reduced SAP HANA license and maintenance costs. We notice that as the volume of data grows, the cost of ownership of SAP HANA doubles from 1.4 million dollars to 2.8 million dollars, which is a 200% increase. However, the data increase on SAP HANA + DT model results in an increase of only 1.52 million dollars, which is a 108% increase. This proves that as business needs grow and the volume of data increases, the more economical option would be to move to SAP HANA with dynamic tiering model.

Table 8: Cost Comparison between SAP HANA and SAP HANA DT with Data Volume Increase

Case when (50% Data Increase)		1
Size (1.5)		0,5
HANA License Cost [Millions] (Cost \$150k per 64GB)		\$3.515
Maintenance [Millions] 22% per year		\$0.770
Total cost (millions)		\$4.285
Savings [SAP HANA 75%/ DT 25%] millions		\$1.07
Savings [SAP HANA 50%/ DT 50%] millions		\$2.14
Savings [SAP HANA 25%/ DT 75%] millions		\$3.21

**All the pricing done for the cost models are based on approximate list prices.

Summary

SAP HANA and dynamic tiering solution on FlashArray//m does not impact the query performance even when 75% of data is fetched from dynamic tiering. The superior read performance of FlashStack ensures that there is no compromise on the query performance when data is moved to dynamic tiering from the core memory. The additional advantage of huge savings of cost makes SAP HANA with dynamic tiering a definite choice.

FlashArray//m Sizing for SAP HANA®



Figure 20: FlashArray //m for SAP HANA

FlashArray//m sizing for SAP HANA® must be done in four phases:

- 1. Throughput sizing
- 2. Capacity sizing
- 3. Data reduction of FlashArray//m
- 4. Dynamic tiering capacity sizing

Throughput Sizing

Throughput sizing depends on number of SAP HANA® nodes; below is the sizing recommendation based on number of SAP HANA® nodes:

Table 9: Throughput sizing recommendation

Scale Out SAP HANA® Nodes	FlashArray//m Controller Model
6	//m20
9	//m50
12	//m70

Capacity Sizing

After throughput sizing is performed, capacity sizing of FlashArray//m is performed, which is based on the overall amount of memory needed by the SAP HANA® database. This is the result of conducting a memory sizing for a given system.

The two types of data (for which memory is required) are:

- **Static data:** For static data, memory is required to hold the database tables. It is calculated by applying a compression factor to the size of the source database and also by performing a data reduction done on FlashArray//m. The value of the compression factor depends on the application scenario such as OLTP (like Suite-on-HANA®) or OLAP (like BW-on-HANA®), and on the source database type (some source databases already apply a certain compression).
- **Dynamic data:** For dynamic data, memory is required to hold objects created dynamically at runtime. For example, at the time of loading new data or when executing queries.

Note: We recommend that you reserve the same amount of memory for dynamic data as well as static data.

FlashArray//m is SAP HANA® Tailored Datacenter Integration (TDI) certified and it can be deployed on hardware infrastructure environments following the SAP HANA® TDI standards. The sizing formula in the following sections are valid for SAP HANA® TDI environments.

Data Volumes Sizing

After a savepoint or a snapshot is created, or a delta merge is performed, data is persisted from memory to the data volume available under /hana/data/<sid>. For the hosts altogether, the recommended size is calculated as follows:



Sizedata = 1 x RAM (= 2 x overall table size)

RAM = Size of memory of SAP HANA® system.

Note: For migration purposes of a non-HANA® database to SAP HANA®, the system may temporarily need more disk space for data than calculated in the sizing phase.

Redo Log Volumes Sizing

The minimum size of the log volume depends on the number of data changes occurring between two SAP HANA® savepoints, which by default are created every 5 minutes. The redo log segments increase with data changes executed by write transactions, and are written to the log volume under /hana/log/<sid>. When sizing the log volume, the following points have to be considered:

- The redo log must not be overwritten before a savepoint entry is available in the data volume; otherwise, the SAP HANA® database may not restart.
- If “log_mode = normal” is set, the redo log must not be overwritten before a backup is performed. Therefore, we recommend having buffer space for incidents or faults that may interrupt the backup process. That buffer space should allow for system administrators to fix and finish the backup process before the log volume fills up.

Below is the formula that calculates the log volume size depending on the RAM:



[systems <= 512GB] Size redo log = 0.5 x RAM (= 1 x overall table size)

[systems > 512GB] Size redo log(min) = 512GB

For high availability configuration of storage replication or SAP HANA® system replication, the storage size of the secondary site must be the same as that of primary.

Shared Volume Sizing

All binary, trace, and configuration files are stored on a shared file system under /hana/shared/<sid>; this is exposed to all hosts of a system. If the host performs a memory dump (which can reach up to 90% of the RAM size), the files are stored in this file system. Thus, additional space is required for the traces written by the compute node(s) of the SAP HANA® database.



SAP HANA® Shared volume = 1 x RAM

Backups

A complete data backup contains the entire payload of all data volumes. The size required by the backup directory not only depends on the total size of the data volumes, but also on the number of backup generations kept on disk and on the frequency with which data is changed in the SAP HANA® database. For example, if the backup policy requires to perform complete data backups on a daily basis and to retain those backups for one week, the size of the backup storage must be seven times the size of the data area.

In addition to data backups, backup storage for log backups must be reserved to provide the possibility for a point-in-time database recovery. The number and size of log backups to be written depend on the number of change operations in the SAP HANA® database.



Size of Backups >= 0.5 x RAM (= 1 x overall table size)

Total Capacity Size

Calculate the total capacity size of SAP HANA® based on above information



SAP HANA® = **Data volumes size + Redo log volumes size**
Total Capacity size + **Shared volume size + Backup size**

Data Reduction of FlashArray//m

Once the throughput calculation and capacity calculation is derived, determine the capacity size by applying the applicable data reduction (refer to data reduction topic in this document under Value propositions section).



Final SAP HANA® Total Capacity size = Data Reduction * SAP HANA® Total Capacity size

Dynamic Tiering Capacity Sizing

Perform sizing of HANA DT which will be based on the following rules. The recommended ratios of HANA memory to DT storage are below:

For HANA memory <= 2.5TB



Size of DT storage should not exceed 4x the size of HANA memory.

Warm data volume = 3 times of SAP HANA memory

DT log volume = 1 time of SAP HANA memory

For HANA memory > 2.5TB:



Size of DT storage should not exceed 8x the size of HANA memory.

Warm data volume = 6 times of SAP HANA memory

DT log volume = 2 times of SAP HANA memory

Conclusions

In this paper, we have demonstrated that dynamic tiering provides the ultimate value in terms of cost per performance to your SAP HANA deployment. We analyzed in detail the read performance for your analytical queries even when they were returning data from dynamic tiering and not from core SAP HANA database. This is primarily due to mighty performance of the Pure Storage all-flash FlashStack solution. A summary of benefits are captured below:

- **Simplicity**—FlashStack eliminates the need for space management of a SAP HANA + dynamic tiering system thereby reducing the operational overhead on the system administrator. Deployment times are reduced via pre-tested and validated designs.
- **Superior read performance**— FlashStack ensures superior read bandwidth which helps with analytical queries read performance even when they are fetching data from dynamic tiering and not from SAP HANA for most of the result set. This reduces the overall TCO by eliminating investment on expensive SAP HANA licensing and certified hardware.
- **Superior write performance**— FlashStack ensures superior write bandwidth and consistent sub-millisecond response time helping SAP HANA® analytical solutions with faster accelerated warm data readiness. This reduces the downtime for the tables which are moved to dynamic tiering.
- **Lower Cost**— FlashStack can dramatically reduce the SAP HANA® with dynamic tiering solutions footprint with best-in-class data reduction over and above the SAP HANA® compression and Sybase IQ compression on which dynamic tiering is based. FlashStack is an ideal platform for consolidating multiple system copies efficiently and economically as data reduction increases even further. Furthermore, a smaller footprint means less power, less datacenter space, and lower cooling requirements, which makes the TCO for large scale SAP HANA with dynamic tiering deployments even more attractive.

Appendix 1: Pure Storage Components in FlashStack

FlashArray //m

The FlashArray//m expands upon the FlashStack's modular, stateless architecture, designed to enable expandability and upgradability for generations. FlashArray//m leverages a chassis-based design with customizable modules, enabling both capacity and performance to be independently improved over time with advances in compute and flash, to meet your business' needs today and tomorrow.

FlashArray//m makes server and workload investments more productive, while also lowering storage spend. With FlashArray//m, organizations can dramatically reduce the complexity of storage to make IT more agile and efficient, accelerating your journey to the cloud.

FlashArray//m's performance can also make your business smarter by unleashing the power of real-time analytics, driving customer loyalty, and creating new, innovative customer experiences that simply weren't possible with disk. All by Transforming Your Storage with FlashArray//m.

FlashArray//m enables you to transform your data center, cloud, or entire business with an affordable all-flash array capable of consolidating and accelerating all your key applications.

Mini Size—Reduce power, space and complexity by 90%

- 3U base chassis with 15-120+ TBs usable
- ~1kW of power
- 6 cables

Mighty Performance—Transform your datacenter, cloud, or entire business

- Up to 300,000 32K IOPS
- Up to 9 GB/s bandwidth
- <1ms average latency

Modular Scale—Scale FlashArray//m inside and outside of the chassis for generations

- Expandable to ~½ PB usable via expansion shelves
- Upgrade controllers and drives to expand performance and/or capacity

Meaningful Simplicity—Appliance-like deployment with worry-free operations

- Plug-and-go deployment that takes minutes, not days
- Non-disruptive upgrades and hot-swap everything
- Less parts = more reliability

The Pure Storage FlashArray//m is ideal for:

Accelerating Databases and Applications—Speed transactions by 10x with consistent low latency, enable online data analytics across wide datasets, and mix production, analytics, dev/test, and backup workloads without fear.

Virtualizing and Consolidating Workloads—Easily accommodate the most IO-hungry Tier 1 workloads, increase consolidation rates (thereby reducing servers), simplify VI administration, and accelerate common administrative tasks.

Delivering the Ultimate Virtual Desktop Experience—Support demanding users with better performance than physical desktops, scale without disruption from pilot to >1000's of users, and experience all-flash performance for under \$100/desktop.

Protecting and Recovering Vital Data Assets—Provide an always-on protection for business-critical data, maintain performance even under failure conditions, and recover instantly with FlashRecover.

Pure Storage FlashArray sets the benchmark for all-flash enterprise storage arrays. It delivers:

Consistent Performance—FlashArray//m delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.

Less Cost than Disk—Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure.

Mission-Critical Resiliency— FlashArray//m delivers >99.999% proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.

Disaster Recovery Built-In— FlashArray//m offers native, fully-integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. And, recover instantly from local, space-efficient snapshots or remote replicas.

Simplicity Built-In— FlashArray//m offers game-changing management simplicity that makes storage installation, configuration, provisioning and migration a snap. No more managing performance, RAID, tiers or caching. Achieve optimal application performance without any tuning at any layer. Manage the FlashArray//m the way you like it: Web-based GUI, CLI, VMware vCenter, Rest API, or OpenStack.




Figure 21: FlashArray//m Specifications

The following table outlines the FlashArray //m specifications for different models:

Table 10: FlashArray//m Specification Matrix

	//m20	//m50	//m70
Capacity	<ul style="list-style-type: none"> Up to 120+ TBs effective capacity* 5 – 40TBs raw capacity (base chassis) 	<ul style="list-style-type: none"> Up to 250+ TBs effective capacity* 30 – 88TBs raw capacity (w/shelves) 	<ul style="list-style-type: none"> Up to 400+ TBs effective capacity* 44 – 136TBs raw capacity (w/shelves)
Performance	<ul style="list-style-type: none"> Up to 150,000 32K IOPS** <1ms average latency Up to 5 GB/s bandwidth 	<ul style="list-style-type: none"> Up to 220,000 32K IOPS** <1ms average latency Up to 7 GB/s bandwidth 	<ul style="list-style-type: none"> Up to 300,000 32K IOPS** <1ms average latency Up to 9 GB/s bandwidth
Connectivity	<ul style="list-style-type: none"> 8 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI Management and Replication ports 	<ul style="list-style-type: none"> 16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI Management and Replication ports 	<ul style="list-style-type: none"> 16 Gb/s Fibre Channel 10 Gb/s Ethernet iSCSI Management and Replication ports
Physical	<ul style="list-style-type: none"> 3U 742 Watts (nominal draw) 110 lbs. (49.9 kg) fully loaded 5.12" x 18.94" x 29.72" FlashArray//m chassis	<ul style="list-style-type: none"> 3U – 7U 1007 - 1447 Watts (nominal draw) 110 lbs. (49.9 kg) fully loaded + 44 lbs. per expansion shelf 5.12" x 18.94" x 29.72" FlashArray//m chassis	<ul style="list-style-type: none"> 5U – 11U 1439 – 2099 Watts (nominal draw) 110 lbs. (49.9 kg) fully loaded + 44 lbs. per expansion shelf 5.12" x 18.94" x 29.72" FlashArray//m chassis

* Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes the benefit of data reduction with always-on inline deduplication, compression, and pattern removal. Average data reduction is calculated at 5-to-1, below the global average of the FlashArray//m user base.

** Why does Pure Storage quote 32K, not 4K IOPS? The industry commonly markets 4K IOPS, but real-world environments are dominated by IO sizes of 32K or larger. FlashArray//m adapts automatically to 512B-32KB IO for superior performance, scalability, and data reduction.

Purity Operating Environment

Purity implements advanced data reduction, storage management and flash management features, and all features of Purity are included in the base cost of the FlashStack.

Storage Software Built for Flash—The FlashCare technology virtualizes the entire pool of flash within the FlashStack, and allows Purity to both extend the life and ensure the maximum performance of consumer-grade MLC flash.

Granular and Adaptive—Purity Core is based upon a 512-byte variable block size metadata layer. This fine-grain metadata enables all of Purity's data and flash management services to operate at the highest efficiency.

Best Data Reduction Available—FlashReduce implements five forms of inline and post-process data reduction to offer the most complete data reduction in the industry. Data reduction operates at a 512-byte aligned variable block size, to enable effective reduction across a wide range of mixed workloads without tuning.

Highly Available and Resilient—FlashProtect implements high availability, dual-parity RAID-3D, non-disruptive upgrades, and encryption, all of which are designed to deliver full performance to the FlashStack during any failure or maintenance event.

Backup and Disaster Recovery Built In—FlashRecover combines space-saving snapshots, replication, and protection policies into an end-to-end data protection and recovery solution that protects data against loss locally and globally. All FlashProtect services are fully-integrated in the FlashStack and leverage the native data reduction capabilities.



Pure1 Manage—By combining local web-based management with cloud-based monitoring, Pure1 Manage allows you to manage your FlashStack wherever you are – with just a web browser.

Pure1 Connect—A rich set of APIs, plugin-is, application connectors, and automation toolkits enable you to connect FlashStack to all your data center and cloud monitoring, management, and orchestration tools.

Pure1 Support—FlashStack is constantly cloud-connected, enabling Pure Storage to deliver the most proactive support experience possible. Highly trained staff combined with big data analytics help resolve problems before they start.

Pure1 Collaborate—Extend your development and support experience online, leveraging the Pure1 Collaborate community to get peer-based support, and to share tips, tricks, and scripts.

Experience Evergreen™ Storage



Tired of the 3-5 year array replacement merry-go-round? The move to FlashStack can be your last data migration. Purchase and deploy storage once and once only – then expand capacity and performance incrementally in conjunction with your business needs and without downtime. Pure Storage's vision for Evergreen Storage is delivered by a combination of the FlashStack's stateless, modular architecture and the ForeverFlash business model, enabling you to extend the lifecycle of storage from 3-5 years to a decade or more.

Appendix 2: Cisco Components in FlashStack

The Cisco Unified Computing System™ (Cisco UCS™) is a next-generation data center platform that unites compute, network, storage access, and virtualization into an organized structure aimed to reduce total cost of ownership and introduce vastly improved infrastructure deployment mechanisms at scale. UCS incorporates a unified network fabric with scalable, modular and powerful x86-architecture servers. With an innovative and proven design, Cisco UCS delivers an architecture that increases cost efficiency, agility, and flexibility beyond what traditional blade and rack-mount servers provide. Cisco makes organizations more effective by addressing the real problems that IT managers and executives face and solves them on a systemic level.

Greater Time-on-Task Efficiency

Automated configuration can change an IT organization's approach from reactive to pro-active. The result is more time for innovation, less time spent on maintenance, and faster response times. These efficiencies allow IT staff more time to address strategic business initiatives. They also enable better quality of life for IT staff, which means higher morale and better staff retention—both critical elements for long-term efficiency.

Cisco UCS Manager is an embedded, model-based management system that allows IT administrators to set a vast range of server configuration policies, from firmware and BIOS settings to network and storage connectivity. Individual servers can be deployed in less time and with fewer steps than in traditional environments. Automation frees staff from tedious, repetitive, time-consuming chores that are often the source of errors that cause downtime, making the entire data center more cost effective.

Easier Scaling

Automation means rapid deployment, reduced opportunity cost, and better capital resource utilization. With Cisco UCS, rack-mount and blade servers can move from the loading dock and into production in a “plug-and-play” operation. Automatically configure blade servers using predefined policies simply by inserting the devices into an open blade chassis slot. Integrate rack-mount servers by connecting them to top-of-rack Cisco Nexus® fabric extenders. Since policies make configuration automated and repeatable, configuring 100 new servers is as straightforward as configuring one server, delivering agile, cost-effective scaling.

Virtual Blade Chassis

With a separate network and separate management for each chassis, traditional blade systems are functionally an accidental architecture based on an approach that compresses all the components of a rack into each and every chassis. Such traditional blade systems are managed with multiple management tools that are combined to give the illusion of convergence for what is ultimately a more labor-intensive, error-prone and costly delivery methodology. Rack-mount servers are not integrated




Figure 22: Cisco Unified Computing System

and must be managed separately or through additional tool sets, adding complexity, overhead, and the burden of more time.

Architecturally, Cisco UCS blade and rack-mount servers are joined into a single virtual blade chassis that is centrally managed yet physically distributed across multiple blade chassis, rack-mount servers, and even racks and rows. This capability is delivered through

Cisco® fabric interconnects that provide redundant connectivity, a common management and networking interface, and enhanced flexibility. This larger virtual chassis, with a single redundant point of management, results in lower infrastructure cost per server, with fewer management touch points, and lower administration, capital, and operational costs.

Reference

SAP HANA® best practices on pure storage guide:

<http://blog.purestorage.com/in-memory-database-meets-mighty-performance-storage-flasharray-is-certified-for-sap-hana/>

Pure Storage FlashArray Data Sheet:

http://www.purestorage.com/content/dam/purestorage/pdf/datasheets/PureStorage_FlashArraym-Brochure.pdf

“HANA Performance Test Benchmarks: <https://scn.sap.com/people/sheetal.jain/blog/2015/05/21/hana-performance-test-benchmarks#>”



Pure Storage, Inc.
Twitter: [@purestorage](https://twitter.com/purestorage)
www.purestorage.com

650 Castro Street, Suite #260
Mountain View, CA 94041

T: 650-290-6088
F: 650-625-9667

Sales: sales@purestorage.com
Support: support@purestorage.com
Media: pr@purestorage.com
General: info@purestorage.com