Pure Storage FlashStack for IoT Solution

A Reference Architecture for Using Virtualized SAP HANA® with Dynamic Tiering on Pure Storage®

White Paper

SAP Co-Innovation Lab

ACKNOWLEDGMENTS

This document is the work of a virtual project team, whose members include:
SAP Co-Innovation Lab: Kevin Liu (project lead), David Cruickshank, Syed Hussaini, Stan Arnbrister
Pure Storage: Krishna Satyavarapu (project lead), John Karagozian

The dedication, hard work, and extra hours that the team has put in during the past three months has made this project successful.

We would also like to thank Courtney Claussen and Sudipto Chowdhuri of the SAP HANA dynamic tiering product team for reviewing and providing expert feedback on this paper.

TABLE OF CONTENTS

Pure Storage FlashStack 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 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 with unique identifiers that can 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 in 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 can also reduce the overall total cost of ownership (TCO) for a Big Data application like IoT. This solves the biggest hurdle enterprises often face when dealing with huge volumes of SAP HANA data.

Based on testing conducted in SAP® Co-Innovation Lab (COIL) Silicon Valley, this paper will analyze the cost and performance advantage that SAP HANA dynamic tiering powered by FlashStack offers over traditional database models. We will also run a comparison with maintenance of data on Oracle to understand the enhanced performance and reduced TCO presented by dynamic tiering on FlashStack.

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 the dynamic tiering model, the need to scale out or scale up SAP HANA storage is eliminated. As the volume of data grows, SAP HANA maintenance costs escalate with licensing and hardware. With FlashStack and dynamic tiering, storage can be expanded with a marginal increase in cost and with an acceptable 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 VMware vSphere 5.5. Pure Storage has validated this reference architecture within the SAP Co-innovation Lab environment. It includes best practices for the benefit of Basis administrators running SAP HANA, storage administrators, and especially for SAP HANA modelers and SAP HANA dynamic tiering modelers.

TARGET AUDIENCE

The target audience for this document includes storage architects, data center architects, Basis administrators running SAP HANA, and especially, SAP HANA modelers and users of SAP HANA dynamic tiering 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 reading this document.

DESIGN PRINCIPLES

The paper explains SAP HANA reference architecture that can be used to manage the ever-compounding volume of data in an IoT scenario. By using an extended store, warm data is moved to a storage device without compromising performance or increasing cost. Data-aging rules can be defined by administrators, and are then automatically executed at scheduled intervals. With the help of reporting tools, data can then be analyzed to present a pattern that enables decision makers to arrive at conclusions. 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 network interface cards (NICs) for connectivity, and high-availability (HA) clustering.
- Efficient: Build a solution that leverages the efficiency 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 the meet the needs of a growing enterprise.

FLASHSTACK INTRODUCTION

FlashStack 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, FlashStack provides the performance and reliability your business applications demand.

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
Consolidation of hundreds of enterprise-class applications in a single rack
Easy scalability, without disruption
Repeatable growth through multiple FlashStack CI deployments

- Operational simplicity

A fully tested, validated, and documented solution for rapid deployment
Reduced management complexity
Auto-aligned 512-byte architecture that eliminates storage alignment headaches
No storage tuning or tiers necessary

- Lowest TCO

Figure 1: Pure Storage FlashStack

Dramatic savings in power, cooling, and space with 100% flash.
Industry-leading data reduction
Free controller upgrades with three-year maintenance renewal under Forever Flash™
Free FlashArray controller upgrades every three years with Forever Flash

- Enterprise-grade resiliency

  Highly available architecture and redundant components
  Nondisruptive operations
  Upgrades and expansion without downtime or performance loss
  Native data protection: snapshots and replication

To learn more about the various components in FlashStack, see Appendix 1: Pure Storage Components in FlashStack and Appendix 2: Cisco Components in FlashStack.

SAP HANA 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 in the cloud. It is a revolutionary platform that is best suited for performing real-time analytics, and developing and deploying real-time applications. As a combination of hardware and software, it is made to process massive real-time data using in-memory computing. (An in-memory database means that 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, 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 that supports relational, graphical, and text data within same system. Its features support significant processing speed, handle huge data volume, and offer text mining capabilities.

Figure 2: SAP HANA® Platform

Dynamic tiering is an add-on option to the SAP HANA database, and is a native data-aging solution to manage larger volumes of less frequently accessed data. Dynamic tiering allows for warm data to be moved to extended storage, which results in reduced size and better performance of the in-memory database.

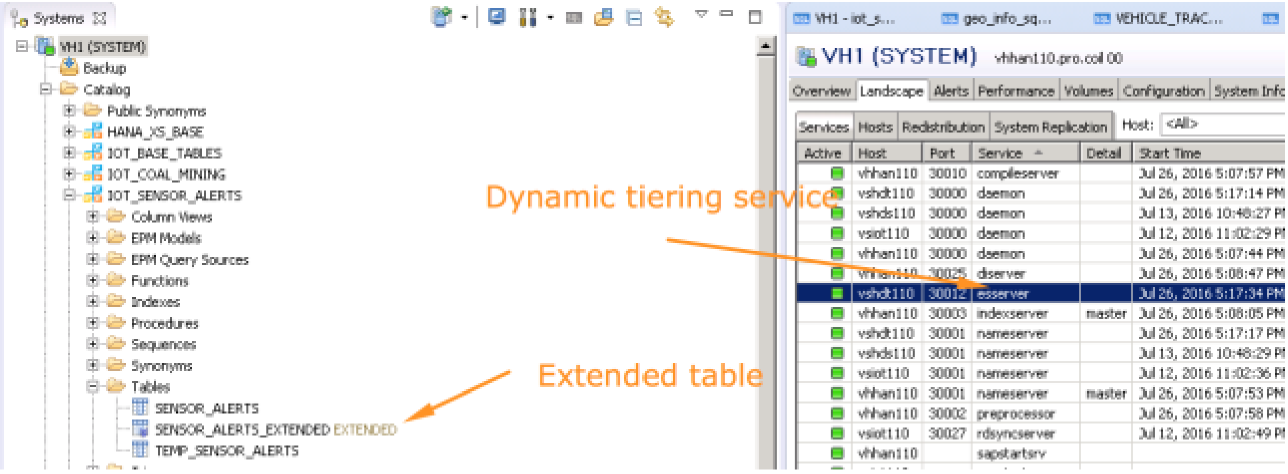

Figure 3: SAP HANA Dynamic Tiering Tables

Dynamic tiering adds settings in the configuration file called esserver.ini 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, the all-flash-based FlashStack. This offloads storage and processing of data that has aged (warm data) from the core SAP HANA memory. The SAP Information Lifecycle Management component is an application powered by SAP HANA 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 lets you schedule those procedures for automatic execution.

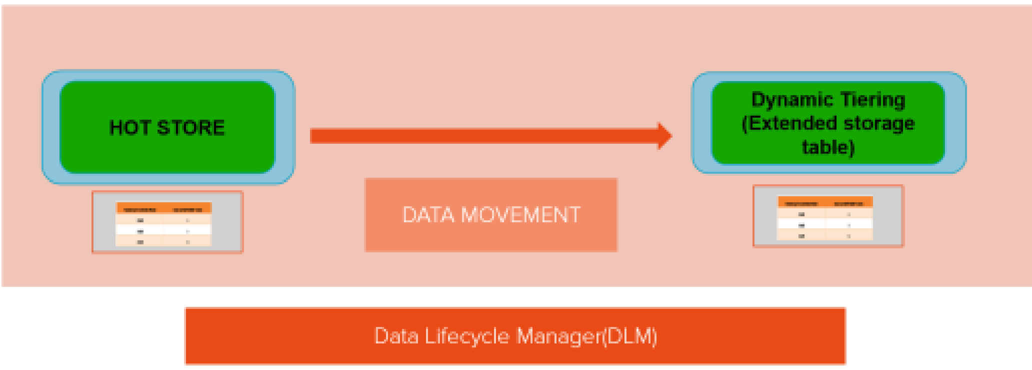

Figure 4: SAP HANA Data Tiering – Data Transfer

The SAP HANA dynamic tiering platform built on Pure Storage all-flash enterprise storage boosts performance of warm data and 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 is not a cost-effective approach. 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.

In the reference architecture we tested at SAP Co-Innovation Lab, SAP BusinessObjects™ Lumira is used as an example of a reporting tool on SAP HANA dynamic tiering. SAP BusinessObjects Lumira is a self-service solution that allows analysts and decision makers to access, transform, and visualize data. The SAP BusinessObjects 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 BusinessObjects Lumira server and cloud platforms (which provide browser- and mobile-based experiences), you can further analyze data and collaborate with colleagues on data sets, stories, and other business intelligence artifacts.

The following use case demonstrates a coal mininge organization leveraging IoT applications to implement a ground-level alert system. Coal mines are complex work environments where safety of workers is of utmost priority; a number of precautions have to be in place to avoid any mishaps. To aid this process, various networked sensors monitor the environment closely, reporting any deviation from an ideal situation. The data is sent to data servers and is monitored regularly with Web applications and servlets. As Figure 5 depicts, alerts are triggered when the system senses smoke, fire, earthquake, equipment failure, and so on. The data is further broken down for each site for each quarter, which can then be analyzed in various ways.

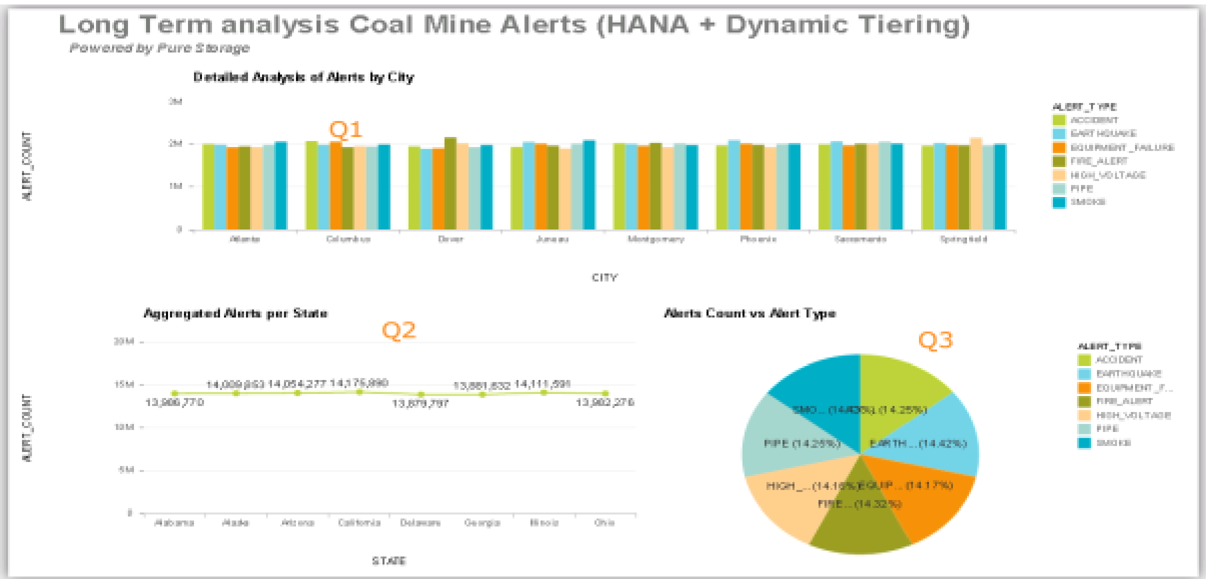

Figure 5: SAP Data Reporting

DEPLOYMENT OF SAP HANA WITH 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 BusinessObjects Design Studio; SAP HANA Cloud Platform, smart data streaming; SAP BusinessObjects Lumira, server version for teams; and Pure Storage GUI). For testing at the lab, we used the following configurations:

- Network: Cisco Unified Computing System (UCS) fabric interconnect 6120XP for external and internal connectivity of IP and fiber channel (FC) network.
- Storage: Pure Storage FlashArray//m20 with fiber channel connectivity
- Compute: Cisco UCS B250 M2 Server

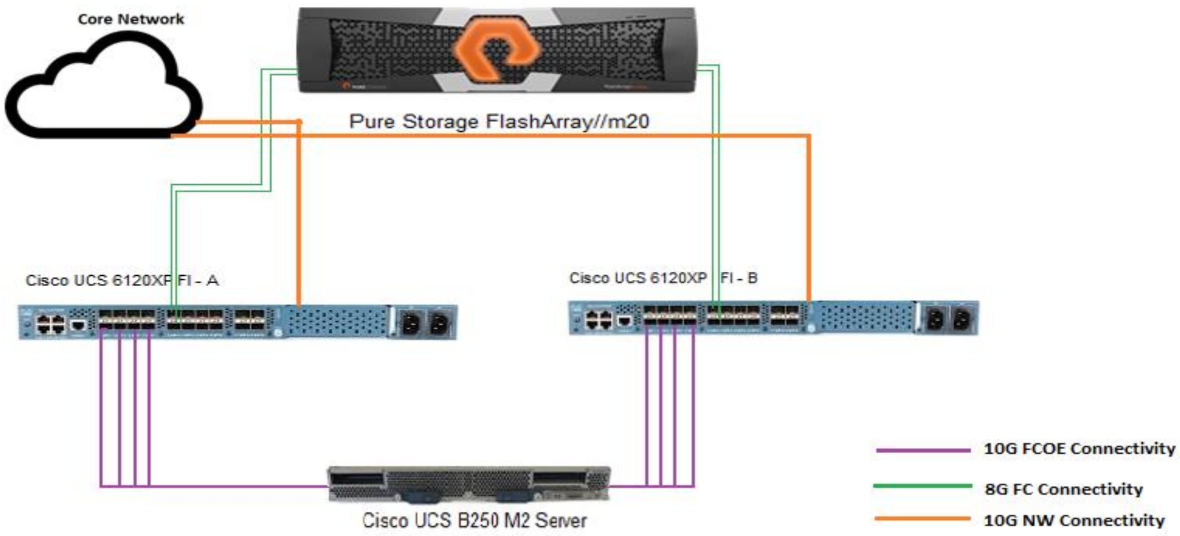

Figure 6: FlashStack Connectivity

Figure 6 shows a detailed topology of the reference architecture configuration deployed at SAP Co-Innovation Lab Silicon Valley. A major goal of the architecture is to build a highly redundant and resilient infrastructure. We used powerful servers with dual FC ports connected redundantly to two fabric interconnect 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 dynamic tiering is completely virtualized using VMware 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 BusinessObjects Design Studio; SAP HANA Cloud Platform, smart data streaming; and SAP BusinessObjects Lumira, server version for teams. 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. 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 logical unit numbers (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
Host bus adapter (HBA)	4 X 16 GB ports on Cisco UCS M81KR/M71KR-Q
Network interface card (NIC)	2 X 10 GB ports on Cisco UCS M81KR/M71KR-Q
Application	SAP HANA dynamic tiering; SAP BusinessObjects Design Studio; SAP HANA Cloud Platform, smart data streaming; SAP BusinessObjects Lumira, server version for teams
Database/platform	SAP HANA SPS11
Virtualization	VMware 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 virtual interface 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 Dynamic Tiering

The setup simulated for this test scenario comprised the below configuration:
A single node, completely virtualized environment using VMware vSphere 5.5.
One Cisco UCS B250 M2 blade used with SUSE 11 SP3 OS with SAP HANA dynamic tiering; SAP BusinessObjects Design Studio; SAP HANA Cloud Platform, smart data streaming; SAP BusinessObjects Lumira, server version for teams, and FlashArray//m20.
Four 16 GB fiber channel ports from the array attached to the SAN Fabric; each node with two fiber 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

Component	Software
Database/platform	SAP HANA SPS11, dynamic tiering, SAP HANA Cloud Platform, smart data streaming
Operating system	SUSE Linux Enterprise 11 Service Pack 3
Reporting application	SAP BusinessObjects Design Studio 16.2.1, SAP BusinessObjects Lumira, server version 1.30
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 SAP HANA , smart data streaming.

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

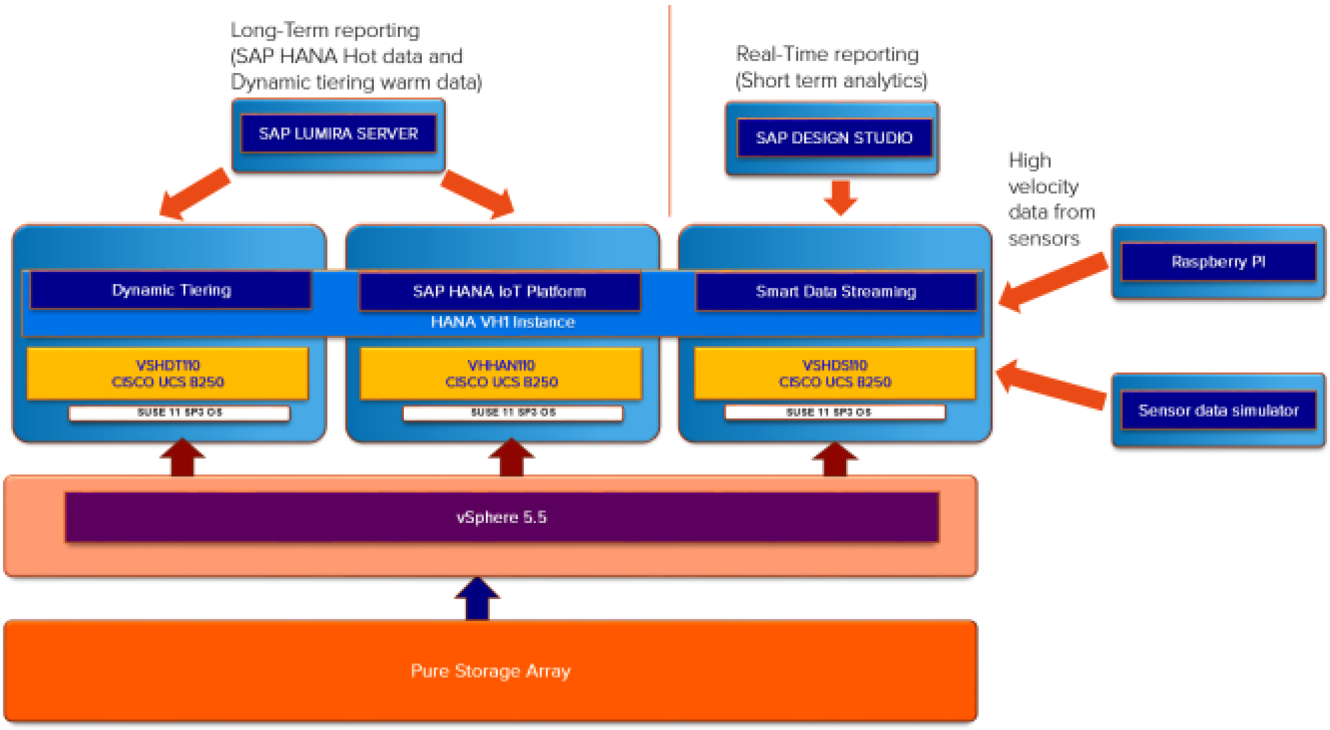

Figure 8: SAP HANA and Dynamic Tiering Setup

We used SAP BusinessObjects 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 SAP Business Explorer (SAP BEx) tools such as the 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 iPad support, a state-of-the art HTML5 UI, built-in application theming, a WYSIWYG Eclipse-based designer, full and native support of SAP BEx queries, direct connectivity to SAP HANA, and an advanced scripting engine.

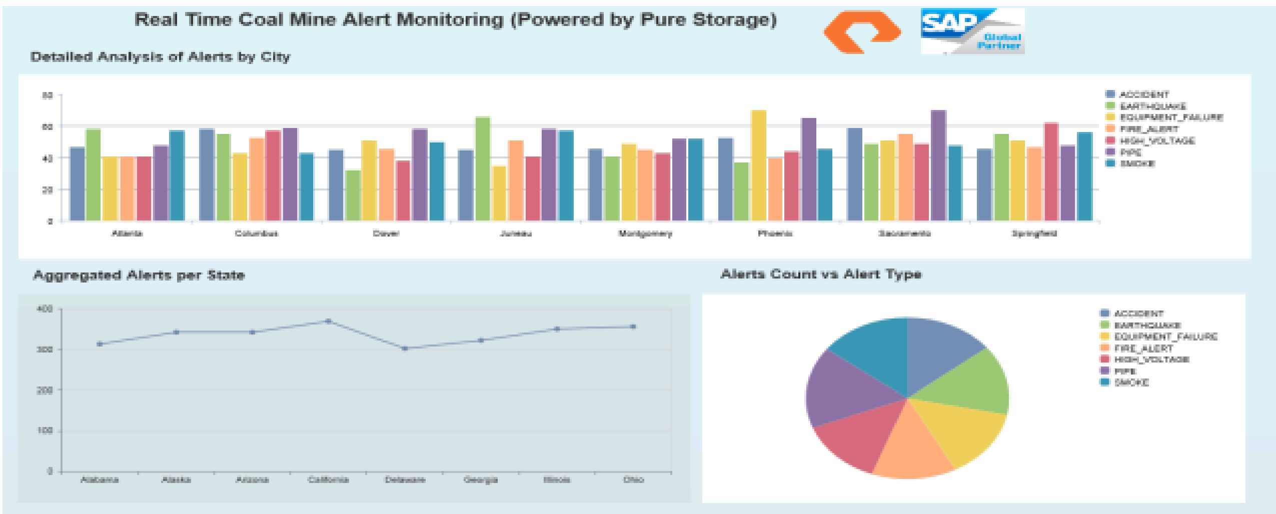

Figure 9: Logical Architecture of the Pure Storage and SAP HANA Setup

The SAP HANA dynamic tiering server must be installed in its own server, which is called an extended storage host. The SAP HANA host and extended storage host or dynamic tiering are required to share the binaries stored in the /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.

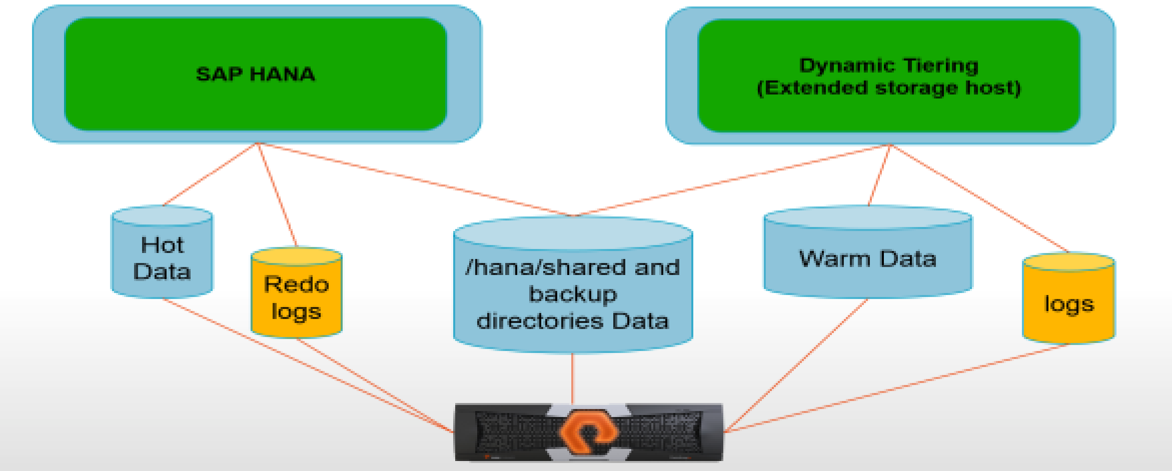

Figure 10: SAP HANA Dynamic Tiering File System

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

Name of the parameter	Description	Recommended settings
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>.
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 250–500 MB for each delta-enabled extended table. This should allow for an adequate safety margin in the event that load frequencies spike on all objects simultaneously.
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 dynamic tiering 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
main_cache_mb	Amount of memory to be used for caching dynamic tiering database objects	Calculate 85% of the total memory on the dynamic tiering 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 continue to explore the value that SAP HANA dynamic tiering on FlashStack offers. The major points of focus are:

- Data reduction
- Accelerated warm data readiness
- Dashboard queries
- Cost and performance analysis

Data Reduction

Dynamic tiering utilizes SAP IQ disk-backed columnar database technology for the extended store. The dynamic tiering store on disk has data compression capabilities that are comparable to those of the SAP HANA in-memory store.

To test data reduction on SAP HANA dynamic tiering, we used 200 GB 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 200 GB 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

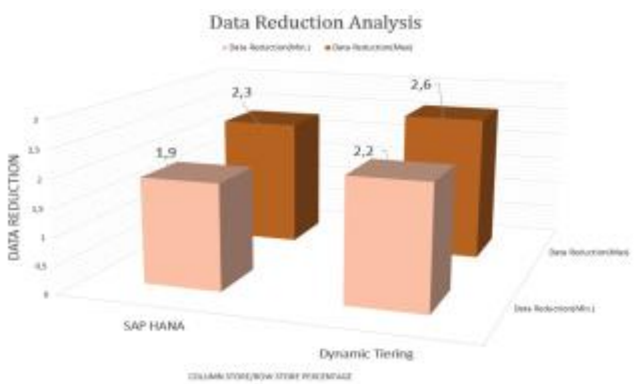

Figure 11: Data Reduction Analysis Graph

Summary

The above results prove that FlashArray//m enhances the data reduction of SAP HANA 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 SAP Information Lifecycle Management 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*>

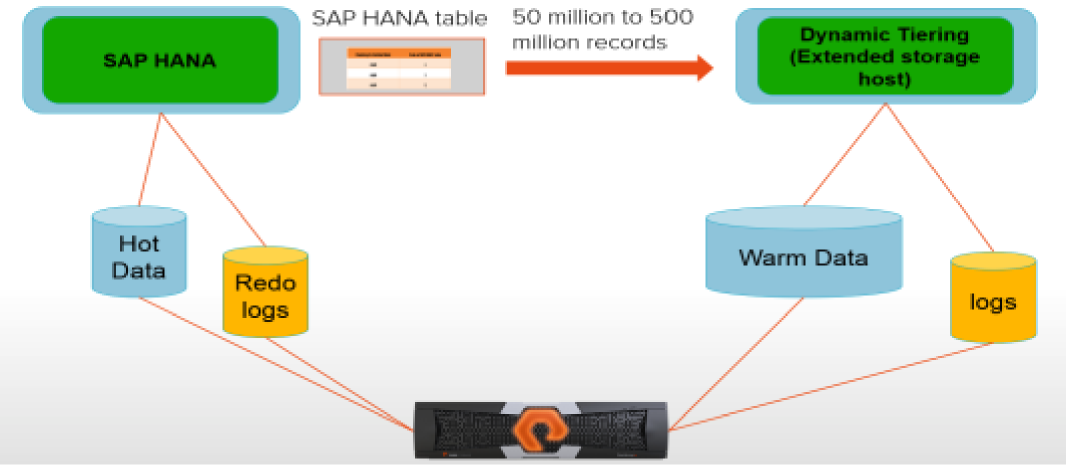

Figure 12: Data Transfer from SAP HANA Using SQL Command

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 dynamic tiering	Peak bandwidth achieved	Avg. bandwidth achieved
50 million	1 minute 3 seconds	63 MB/second	40 MB/second
250 million	4 minutes 17 seconds	278 MB/second	165 MB/second
500 million	7 minutes 23 seconds	385 MB/second	229 MB/second

The following graph depicts 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 the entire data move to dynamic tiering took little more than 4 minutes.

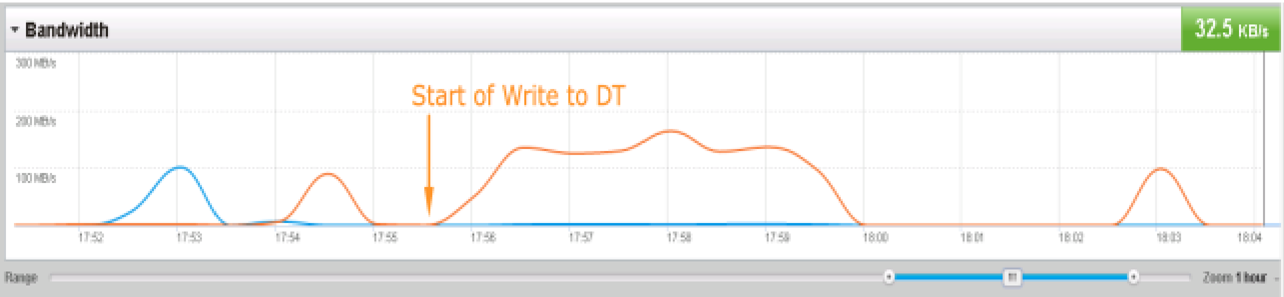

Figure 13: Transfer of 250 Million Records from SAP HANA

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

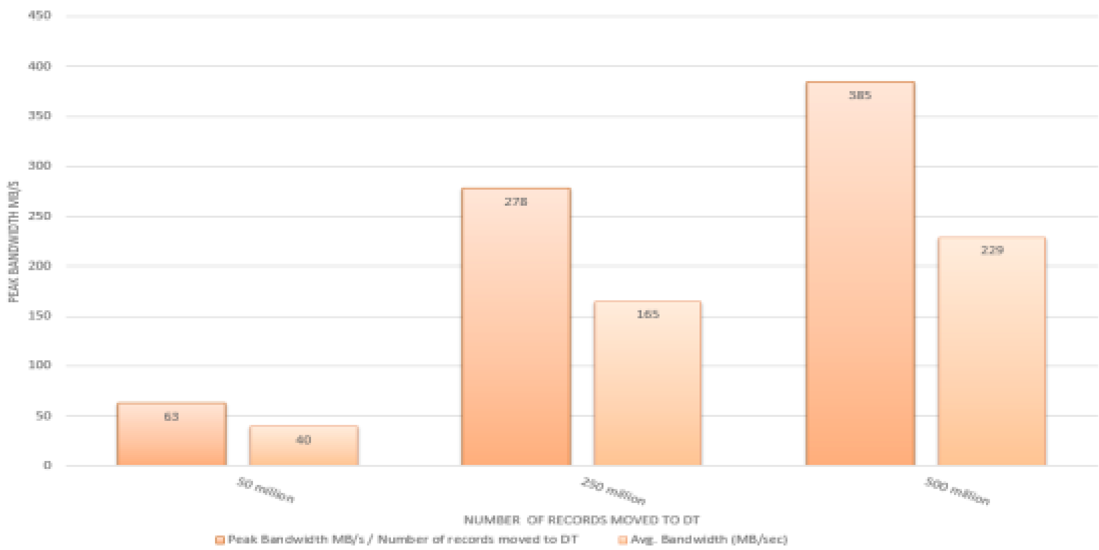

Figure 14: Peak/Average Bandwidth During Data Transfer

Summary

Due to the superior write performance of FlashStack, 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, neither a cost-effective option. 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. Dynamic tiering also offers a TCO advantage compared to SAP HANA, which involves licensing and hardware.

Below is the detailed analysis of how we tested this scenario and a 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 get 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 a 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 >1 billion records. We then scaled down the number of records to zero by pushing the data to the dynamic tiering table, and scaled dynamic tiering up from 0 to 500 million records. We ran the three queries as shown in the SAP BusinessObjects Lumira dashboard below, with all of them doing a 500-million-record scan.

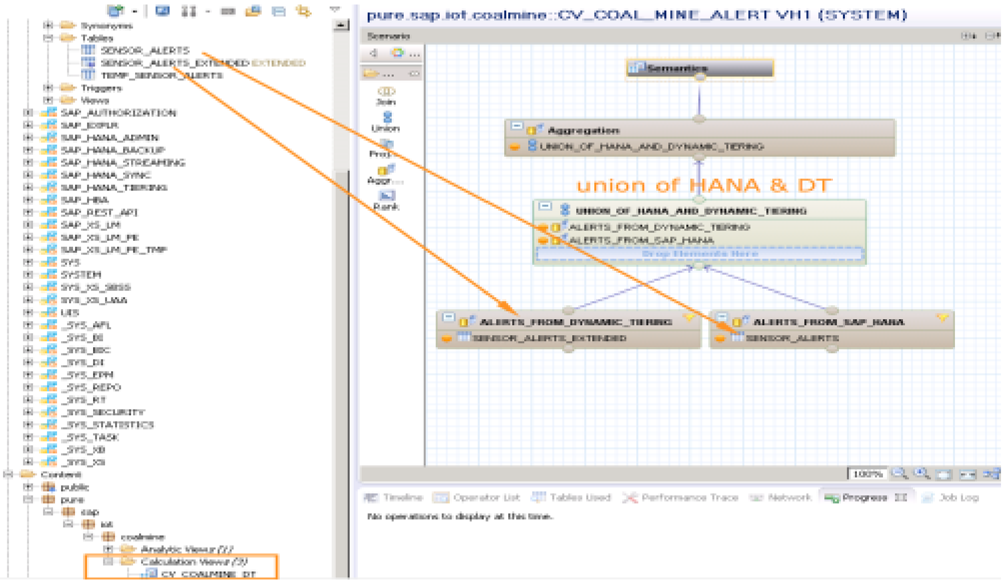

Figure 15: Calculation View and Union of SAP HANA and Dynamic Tiering Tables

Below is the breakdown of SAP BusinessObjects Lumira dashboard queries hitting SAP HANA and the 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.

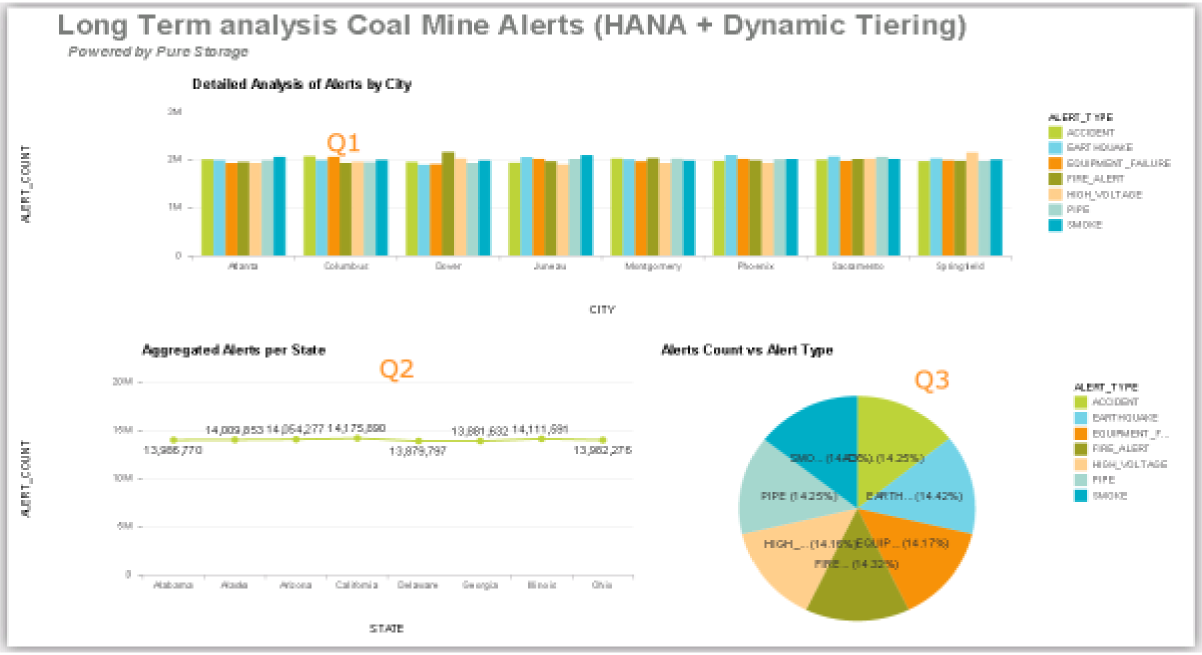

Figure 16: SAP BusinessObjects Lumira Dashboard Queries on Calculation View

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

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

Table 6: IoT Queries

Case 1—100% of data in SAP HANA / 0% of data in dynamic tiering (baseline test case)

These queries were run when all the data was present in SAP HANA and no data present in dynamic tiering.

Query name	Query description	Number of records in SAP HANA	Number of records in dynamic tiering	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% of data in SAP HANA / 25% of data in dynamic tiering

These queries were run when 25% of the data was moved to dynamic tiering, which is around 125 million records.

Query name	Query description	Number of records in SAP HANA	Number of records in dynamic tiering	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	1.318 second
Q3	Alert count vs. alert type	~ 375 million records	~125 million records	1.172 second

Case 3—50% of data in SAP HANA / 50% of data in dynamic tiering

These queries were run when 50% of the data was moved to dynamic tiering, which is around 250 million records.

Query name	Query description	Number of records in SAP HANA	Number of decords in dynamic tiering	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

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 dynamic tiering. We notice that the query runtimes do not vary sharply when compared to baseline runtime when data increases or decreases in dynamic tiering. This shows the superior read performance of the FlashArray//m.

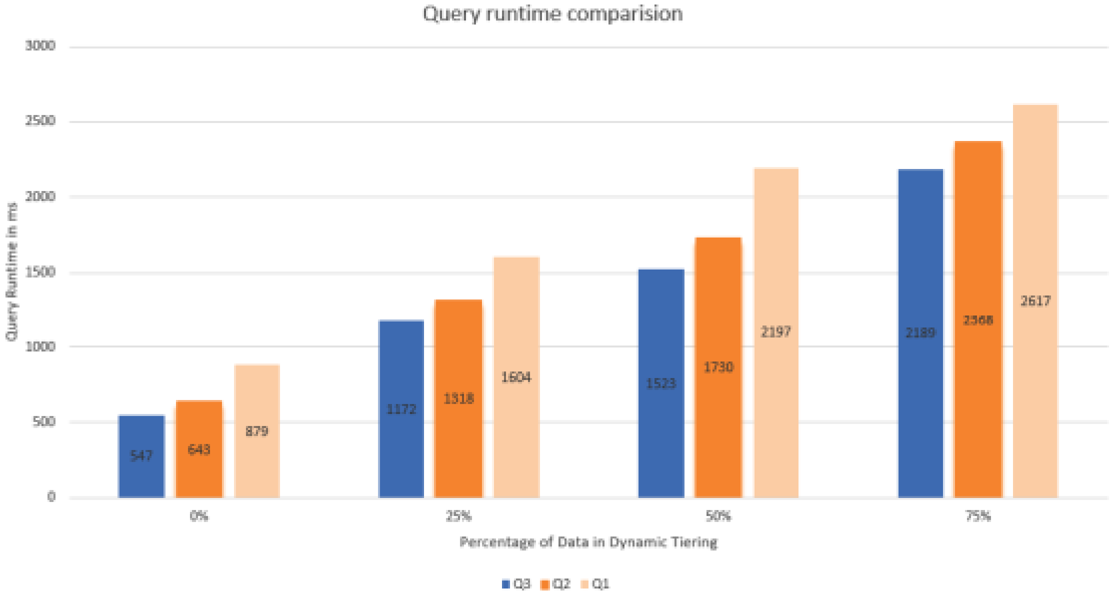

Figure 17: Query Runtime Comparison Graph

Summary

We find that to tackle the scenario of increasing data volumes, increasing dynamic tiering with Pure Storage makes business sense to scale SAP HANA so as not to escalate the TCO.

Cost and Performance Analysis

Let us compare the cost associated with maintenance of data, both in SAP HANA and in dynamic tiering (both running FlashStack) 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 an Oracle database and on SAP HANA with dynamic tiering. The performance analysis on Oracle database is based on the 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 dynamic tiering increments by 25%

In this test scenario, we will compare the performance of query runtime on SAP HANA dynamic tiering with that of Oracle. We tested the cost and performance advantage by reducing the data on SAP HANA by 25% and increasing the dynamic tiering load correspondingly. We will then take a combined view of the query runtime and cost advantage of the data on SAP HANA dynamic tiering with that of data on the Oracle database (based on findings from the SAP HANA public site[1]). Just for the purpose of this comparison the license cost of SAP HANA is $150,000 per 64 GB of data (per the references cited[2]).

[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

SAP HANA with dynamic tiering cost and performance advantage compared to Oracle

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

Table 7: Cost and Performance Comparison (SAP HANA vs. Oracle)

Case	1	2	3	4
Initial SAP HANA size	1TB			
SAP 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 (SAP HANA vs. Oracle)	13.90	11.93	11.26	10.42
SAP HANA license cost [millions] $150K per 64 GB	$ 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 in millions	$ 3.030	$ 2.273	$ 1.515	$ 0.757
Savings in millions	$ -	$ 0.757	$ 1.515	$ 2.273

Figure 18 shows the performance difference between Oracle and SAP HANA with dynamic tiering for queries. This shows that queries run 6–13.9 times faster on SAP HANA with dynamic tiering compared to Oracle. Besides this, there is a substantial savings of spending per year in terms of SAP HANA license cost and hardware maintenance.

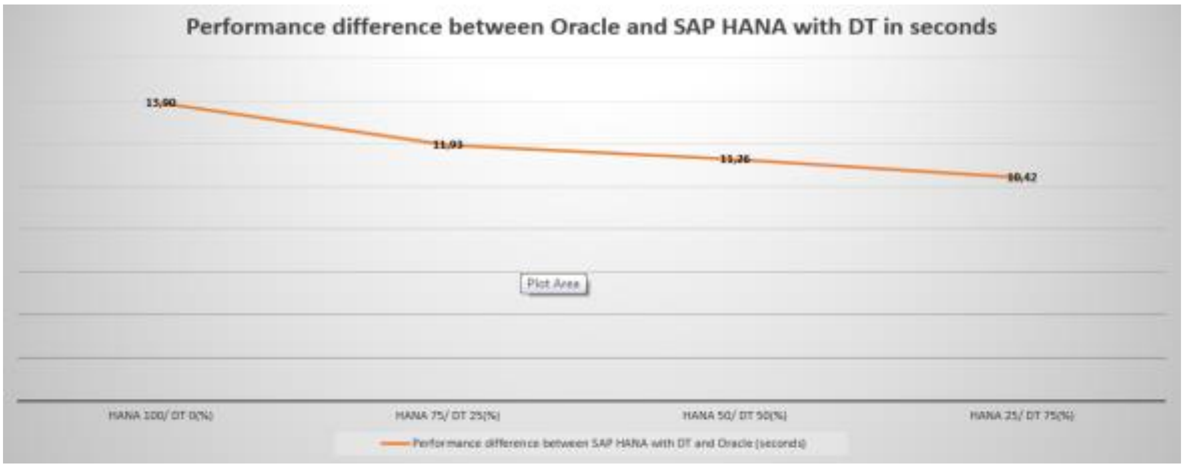

Figure 18: Performance Comparison between SAP HANA with Dynamic Tiering and Oracle

SAP HANA with Dynamic Tiering TCO Advantage as Compared to SAP HANA Scale-Up Strategy

Figure 19 illustrates the economy that the SAP HANA dynamic tiering model offers when compared with scale-up model of SAP HANA, based on 1 TB size. This accounts for cost savings in terms of SAP HANA licensing and hardware maintenance, as more data is moved to dynamic tiering by reducing the SAP HANA footprint.

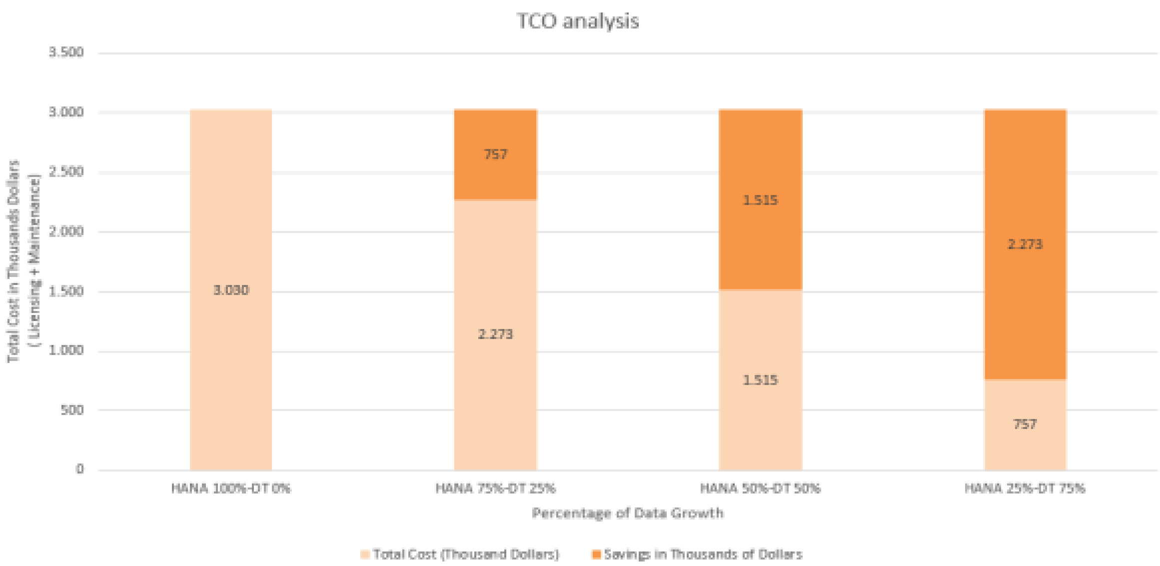

Figure 19: Cost Comparison Between SAP HANA with Dynamic Tiering and SAP HANA Scale-Up

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 of keeping 100% of the data in SAP HANA and of moving it to dynamic tiering. You will notice that even moving 25% of the data to dynamic tiering from SAP HANA, the savings is as much as $1 million per year by way of reduced SAP HANA license and hardware maintenance costs. We notice that as the volume of data grows, the cost of ownership of SAP HANA doubles from $1.4 million to $2.8 million, a 200% increase. However, the data increase on the SAP HANA dynamic tiering model results in an increase of only $1.52 million, or 108%. 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.

Table 8: Cost Comparison Between SAP HANA and SAP HANA Dynamic Tiering with Data Volume Increase

Case with (50% data increase)	1
Size (1.5)	0,5
SAP HANA license cost [millions] (Cost $150K per 64 GB)	$3.515
Maintenance [millions] 22% per year	$0.770
Total cost (millions)	$4.285
Savings [SAP HANA 75%/dynamic tiering 25% – millions]	$1.07
Savings [SAP HANA 50%/dynamic tiering 50% – millions]	$2.14
Savings [SAP HANA 25%/dynamic tiering 75% – millions]	$3.21

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

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

FLASHARRAY//M SIZING 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

Figure 20: FlashArray//m for SAP HANA

Throughput Sizing

Throughput sizing depends on the number of SAP HANA nodes; below is the sizing recommendation based on the 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 SAP Business Suite powered by SAP HANA) or OLAP (like SAP Business Warehouse powered by SAP 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 certified for SAP HANA tailored data center integration, and can be deployed on hardware infrastructure environments following the SAP HANA tailored data center integration standards. The sizing formula in the following sections is valid for SAP HANA tailored data center integration 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 the SAP HANA database

Note: For migration purposes of a non-SAP 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 five 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 should 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 the 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. Note that a complete data backup includes not just the SAP HANA memory store, but the dynamic tiering store as well. 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 performing complete data backups on a daily basis and retaining 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, storage for log backups must be reserved to provide for the possibility of 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 backup >= 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 and capacity calculations are derived, determine the capacity size by applying the applicable data reduction. (Refer to data reduction topic in this document under the “Value Propositions” section above.)

Final SAP HANA total capacity size = Data reduction * SAP HANA total capacity size

Dynamic Tiering Capacity Sizing

The dynamic tiering store achieves comparable data compression rates to SAP HANA. A rough estimate is to set up a total amount of dynamic tiering storage equal to half of the raw data volume size of cool/warm data. For example, if you have 2 TB of raw dynamic tiering data, configure storage capacity at 1 TB for that data. Ensure that the storage can provide at least 50–100 MB per second per core of throughput to the dynamic tiering server.

There is a maximum amount of dynamic tiering storage you may set up compared to SAP HANA memory. The recommended ratios of SAP HANA memory to dynamic tiering storage are below:

For SAP HANA memory <= 2.5 TB

Size of dynamic tiering storage should not exceed 4x the size of SAP HANA memory.
Warm data volume = 3 times of SAP HANA memory
Dynamic tiering log volume = 1 time of SAP HANA memory

For SAP HANA memory > 2.5 TB:

Size of dynamic tiering storage should not exceed 8x the size of SAP HANA memory.
Warm data volume = 6 times of SAP HANA memory
Dynamic tiering log volume = 2 times of SAP HANA memory

Please refer to the following guide for all aspects of sizing the dynamic tiering server (cores, memory, and storage capacity): http://go.sap.com/docs/download/2016/02/8c7bc509-617c-0010-82c7-eda71af511fa.pdf

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 analytical queries even when they were returning data from dynamic tiering and not from the core SAP HANA database. This is primarily due to the powerful performance of Pure Storage all-flash FlashStack. A summary of benefits is captured below:

- Simplicity – FlashStack eliminates the need for space management of a system running SAP HANA with dynamic tiering, thereby reducing the operational overhead on the system administrator. Deployment times are reduced via pretested and validated designs.
- Superior read performance – FlashStack delivers 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 SAP HANA licensing and certified hardware.
- Superior write performance – FlashStack delivers superior write bandwidth and consistent sub-millisecond response time, helping SAP HANA analytical solutions with accelerated warm data readiness. This reduces the downtime for tables moved to dynamic tiering.
- Lower cost – FlashStack can dramatically reduce the footprint of SAP HANA with dynamic tiering with best-in-class data reduction over and above the SAP HANA compression and SAP 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 data center space, and lower cooling requirements, which makes the TCO for large-scale deployments of SAP HANA with dynamic tiering even more attractive.

APPENDIX 1: PURE STORAGE COMPONENTS IN FLASHSTACK

FlashStack comprises a set of hardware and software to perform storage, network, and compute operations. This appendix summarizes the 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, your organization 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
- ~1 kW of power
- 6 cables

Mighty performance – Transform your data center, 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
- Nondisruptive 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 data sets, 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 >1,000'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 real-world application 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 nondisruptively without performance impact.
- Disaster recovery built-in – FlashArray//m offers native, fully integrated, data reduction–optimized backup and disaster recovery at no additional cost. Set up 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.

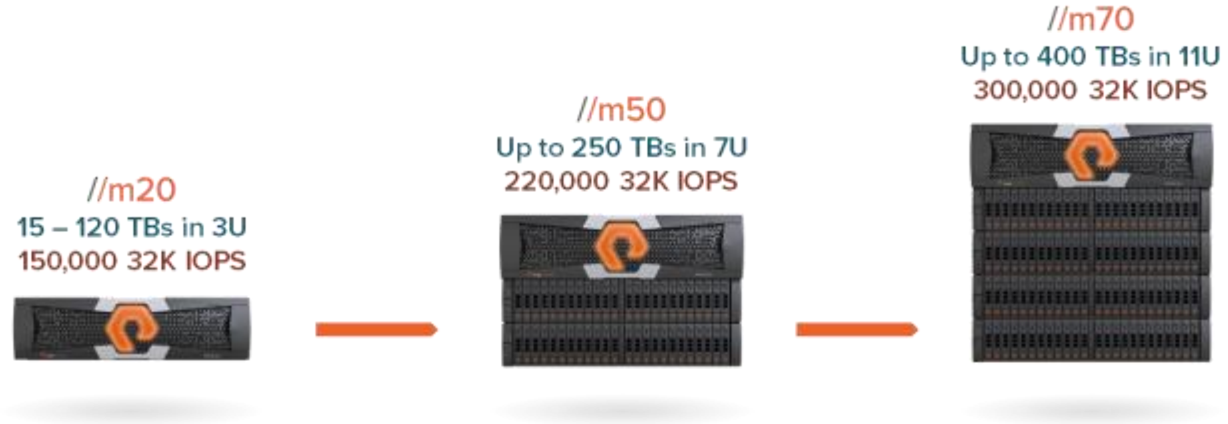

Figure 21: FlashArray//m Specifications

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

	//m20	//m50	//m70
Capacity	Up to 120+ TBs effective capacity* 5–40 TBs raw capacity (base chassis)	Up to 250+ TBs effective capacity* 30–88 TBs raw capacity (w/shelves)	Up to 400+ TBs effective capacity* 44–136 TBs raw capacity (w/shelves)
Performance	• Up to 150,000 32K IOPS** <1ms average latency Up to 5 GB/s bandwidth	• Up to 220,000 32K IOPS** <1ms average latency Up to 7 GB/s bandwidth	• Up to 300,000 32K IOPS** <1ms average latency Up to 9 GB/s bandwidth
Connectivity	8 GB/s fiber channel 10 GB/s Ethernet iSCSI Management and Replication ports	16 GB/s fiber channel 10 GB/s Ethernet iSCSI Management and Replication ports	16 GB/s fiber channel 10 GB/s Ethernet iSCSI Management and Replication ports
Physical	3U 742 Watts (nominal draw) 110 lbs. (49.9 kg) fully loaded • 5.12" x 18.94" x 29.72" FlashArray//m chassis	• 3U – 7U 1,007–1,447 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	• 5U – 11U 1,439–2,099 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

Table 10: FlashArray//m Specification Matrix

* 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 512-byte-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 support 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; nondisruptive 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 its native data-reduction capabilities.

Pure1

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 array replacement merry-go-round every three to five years? 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 three to five 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.

Figure 22: Cisco Unified Computing System

Greater Time-on-Task Efficiency

Automated configuration can change an IT organization's approach from reactive to proactive. 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 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-flasharraym-is-certified-for-sap-hana/

- Pure Storage FlashArray Data Sheet: http://www.purestorage.com/content/dam/purestorage/pdf/datasheets/PureStorage_FlashArraym-Brochure.pdf

www.sap.com