

WHITE PAPER

# Archive and Backup in One for In-Memory and High-Performance Databases

Reference Architecture Guide

By Milind Pathak and Archana Kuppuswamy

October 2017

# Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@HitachiVantara.com](mailto:SolutionLab@HitachiVantara.com). To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

## Revision History

Revision	Changes	Date
SL-008-00	Initial release	October 23, 2017

# Table of Contents

<b>Solution Overview</b>	<b>2</b>
<b>Key Solution Elements</b>	<b>4</b>
Hardware Elements	4
Software Elements	7
<b>Solution Design</b>	<b>7</b>
SAP HANA and S/4HANA Backup and Archival Infrastructure	8
Hitachi Content Platform (HCP)	9
Veritas NetBackup (NBU) Architecture	10
PBS ContentLink Architecture	10
Fibre Channel SAN Architecture	10
Network Architecture	11
Backup and Recovery Operations	13
Archive Operations	14
<b>Engineering Validation</b>	<b>16</b>
Test Methodology	16

# Archive and Backup in One for In-Memory and High-Performance Databases

## Reference Architecture Guide

This reference architecture guide describes the backup, recovery, and archiving of SAP systems and SAP HANA databases in a scale-up configuration using Hitachi Content Platform (HCP).

SAP HANA is an in-memory database that keeps the bulk of its data in memory and uses persistent storage to provide a fallback in case of failure. However, if the persistent storage itself is damaged, data backups are required for protection against data loss due to disk failures or database corruption. The loss of business-critical system resources and services, like the SAP HANA platform, translate directly into lost revenue. Therefore, it is critical to perform SAP HANA database backups to a secondary storage media like external storage or tape libraries, and to be able to restore the SAP HANA database from the backup data in case of a failure.

Continuous usage of SAP enterprise systems results in massive amounts of data. Large databases increase storage cost and impact on database performance, backup window and migration downtime. Data Archiving, a service provided by SAP NetWeaver, provides a solution by removing mass data that the system no longer needs online, but must still be accessible at later date if required, from the database.

In this solution, Hitachi Content Platform (HCP) is used to solve these problems.

- For backup and recovery operations of SAP HANA Platform in a scale-up configuration, this solution uses VERITAS NetBackup (NBU) which is a SAP certified third party backup tool. Using Backint for SAP HANA, NBU can be used to back up the data and log volumes of the SAP HANA system. NBU can use HCP as on-premise cloud object store for storing SAP HANA database backups.
- For archiving the S/4HANA application, this solution uses PBS ContentLink software. PBS Content Link allows customers to archive SAP files and documents, print lists, and provides a lean and very effective solution that links customer's SAP environments. PBS ContentLink can integrate with HCP to be used as an archive destination for S/4HANA.

This reference architecture solution is deployed using the following components:

- Hitachi Compute Blade 2500
- Hitachi Content Platform (HCP)
- Virtual Storage Platform (VSP) G600
- Hitachi Content Platform
- SAP High-Performance Analytic Appliance (HANA)
- VERITAS NetBackup
- PBS ContentLink

Use this document to understand the example reference architecture for archiving, backup, and recovery of SAP HANA in a scale-up configuration, and to help with deployment of the configuration.

This solution was tested in the lab using a 2-socket scale-up configuration only; however, any scale-up Hitachi Unified Compute Platform for SAP HANA configuration is supported with this reference architecture. For information concerning your implementation, please contact your GSS representative for more details.

This technical paper assumes familiarity with the following areas:

- Storage area network-based storage systems
- General storage concepts
- General network knowledge
- SAP HANA platform
- Common IT best practices
- SAP High-Performance Analytic Appliance (HANA)

---

**Note** — These procedures were developed in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before implementing this solution in your production environment. Test the implementation in a non-production, isolated test environment that otherwise matches your production environment.

---

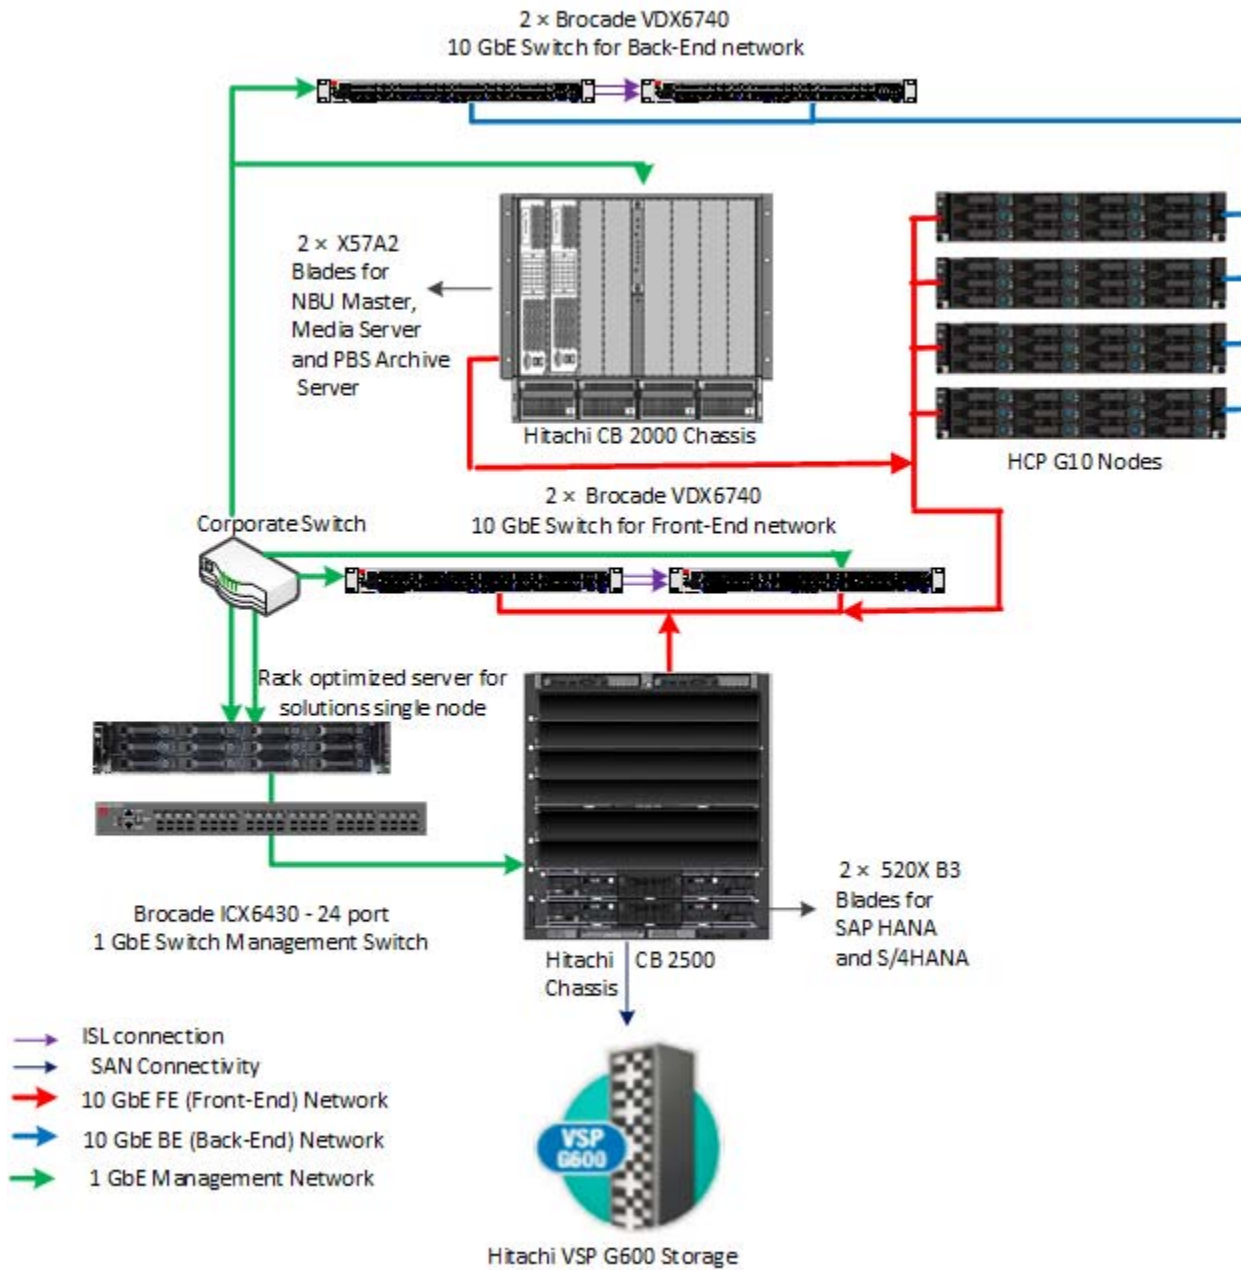
## Solution Overview

This scale-up Hitachi Unified Compute Platform for SAP HANA backup and recovery reference solution using Backint for SAP HANA with NetBackup uses the following components:

- **Hitachi Compute Blade 2500** — An enterprise-class blade server platform. This solution uses two 520X B3 server blades.
- **Rack optimized server for solutions, 2U single node** — A midrange rack mountable server platform, providing advanced systems management and redundancy options. This solution uses rack optimized server for solutions, 2U single node server as a central device for managing all the hardware components used in this solution.
- **Virtual Storage Platform (VSP) G600** — A midrange storage solution. It is used as a primary storage for SAP application and SAP HANA database.
- **Hitachi Content Platform** — Used for handling backup, restore, and archive data storage needs for SAP HANA database and S/4HANA application.
- **SAP HANA platform** — A multi-purpose, in-memory database appliance to analyze transactional and analytical data.
- **Brocade VDX 6740-48 switch** — A 48-port switch that provides 10 GbE connectivity to the appliance. This solution uses four VDX 6740-48 switches.
- **Brocade ICX 6430-48 switch** — A 48-port 1 GbE switch that provides management network to the appliance.
- **VERITAS NetBackup** — Installed on the two X57A2 server blades with Windows servers and acts as NetBackup media and master server. Performs fast and efficient backup of SAP HANA database to the Hitachi Content Platform.
- **PBS Content Link** — PBS ContentLink creates a direct and lean connection between SAP systems and Hitachi Content Platform. This enables archiving that meets legal requirements and a future-proof Information Lifecycle Management concept.

Figure 1 shows the configuration of this solution.

**Figure 1**



## Key Solution Elements

These are the key hardware and software components used in this reference architecture.

### Hardware Elements

Table 1 describes the hardware used to deploy the archive and backup solution for a scale-up node and configuration.

TABLE 1. HARDWARE ELEMENTS

Hardware	Quantity	Configuration	Role
Hitachi Compute Blade 2500 chassis	1	<ul style="list-style-type: none"> <li>■ 8-server blade chassis</li> <li>■ 2 management modules</li> <li>■ 10 cooling fan modules</li> <li>■ Power supply modules                             <ul style="list-style-type: none"> <li>■ 2 modules for 2-socket</li> <li>■ 3 modules for 4-socket</li> <li>■ 4 modules for 8-socket</li> </ul> </li> <li>■ 4 I/O board modules</li> <li>■ 2 × 10 GbE, 2-port LAN PCIe adapter</li> <li>■ 2 × Hitachi 16 Gb/sec 2-port Fibre Channel adapters</li> </ul>	Server blade chassis
520X B3 server blade	2	<ul style="list-style-type: none"> <li>■ 2 × 22-core processors</li> <li>■ 768 GB RAM</li> <li>■ 1 × 2-port pass-through mezzanine card on mezzanine slot 2 and mezzanine slot 4 of server blade 1 and server blade 2</li> </ul>	SAP HANA server
Hitachi Compute Blade 2000 chassis		<ul style="list-style-type: none"> <li>■ 8-blade chassis</li> <li>■ 2 Management modules</li> <li>■ 8 Cooling fan modules</li> <li>■ 4 Power supply modules</li> <li>■ 4 × 1 Gb/sec pass-through modules</li> </ul>	Server blade chassis

TABLE 1. HARDWARE ELEMENTS (CONTINUED)

Hardware	Quantity	Configuration	Role
Hitachi Compute Blade 2000 X57A2 server blade		<ul style="list-style-type: none"> <li>■ 2 × 10 core processors</li> <li>■ 256 GB memory</li> <li>■ 1 × 4 port 1 Gb Ethernet mezzanine card in slot 0</li> </ul>	Veritas NBU Media Server, Master Server and PBS Archive Server
Hitachi Virtual Storage Platform G600	1	<ul style="list-style-type: none"> <li>■ Single frame</li> </ul>	Block storage for SAP HANA nodes and applications
VSP G600 management server (MiniMapp server)	1	Single server	Management server for the storage VSP G600
Rack optimized server for solutions single node	1	<p>1 server node</p> <ul style="list-style-type: none"> <li>■ Intel Xeon E5-2620 v3 processor, 2.4 GHz CPU, 32 GB RAM</li> <li>■ 2 × 500 GB 7200 RPM SATA drives</li> <li>■ 1 dual port 10 GigE Intel 82599ES SFP+OCP Mezzanine card</li> <li>■ 1 dual port 1 GigE Base-T Intel i350 Mezzanine Card</li> <li>■ Emulex Dual Port 8 Gb/sec Fibre Channel HBA</li> </ul>	<ul style="list-style-type: none"> <li>■ SAP HANA Management server</li> </ul>
Hitachi Content Platform (HCP G10)	1	<ul style="list-style-type: none"> <li>■ 4 HCP G10 Nodes</li> <li>■ 2 × Intel E5-2620v3</li> <li>■ 6-core CPUs</li> <li>■ 4 × 16 GB 2133 MHz DDR4 DIMMs per node</li> <li>■ LSI 3108 RAID mezzanine card</li> <li>■ 6 × 4 TB hard disk drives</li> <li>■ 2 × 800 GB 2.5-inch SSDs</li> <li>■ 2 × dual-port 10 Gb Ethernet SFP+ PCIe card</li> </ul>	<ul style="list-style-type: none"> <li>■ Hitachi On-Premise Object Store</li> </ul>



TABLE 1. HARDWARE ELEMENTS (CONTINUED)

Hardware	Quantity	Configuration	Role
Brocade VDX 6740-48 port switch	4	<ul style="list-style-type: none"> <li>■ Two switches for HCP back-end network and two switches for front-end network connecting HCP with rest of the applications</li> </ul>	10 GbE front-end and back-end network
Brocade ICX 6430-48 port switch	1	<ul style="list-style-type: none"> <li>■ 1 GbE</li> <li>■ 48 ports</li> </ul>	1 GbE Management Network

### *Hitachi Compute Blade 2500*

[Hitachi Compute Blade 2500](#) delivers enterprise computing power and performance with unprecedented scalability and configuration flexibility. Lower your costs and protect your investment.

Flexible I/O architecture and logical partitioning allow configurations to match application needs exactly with Hitachi Compute Blade 2500. Multiple applications easily and securely co-exist in the same chassis.

Add server management and system monitoring at no cost with Hitachi Compute Systems Manager. Seamlessly integrate with Hitachi Command Suite in Hitachi storage environments.

### *Hitachi Virtual Storage Platform Gx00 Models*

[Hitachi Virtual Storage Platform Gx00 models](#) are based on industry-leading enterprise storage technology. With flash-optimized performance, these systems provide advanced capabilities previously available only in high-end storage arrays. With the Virtual Storage Platform Gx00 models, you can build a high performance, software-defined infrastructure to transform data into valuable information.

Hitachi Storage Virtualization Operating System provides storage virtualization, high availability, superior performance, and advanced data protection for all Virtual Storage Platform Gx00 models. This proven, mature software provides common features to consolidate assets, reclaim space, extend life, and reduce migration effort. New management software improves ease of use to save time and reduce complexity. The infrastructure of Storage Virtualization Operating System creates a management framework for improved IT response to business demands.

### *Hitachi Content Platform*

[Hitachi Content Platform](#) (HCP) provides distributed object storage for advanced unstructured data storage management. This helps you address challenges of ever-growing volumes of unstructured file storage. Divide a single Content Platform into multiple virtual object stores, secure access to each store, and uniquely configure each store for a particular workload.

Eliminate storage silos using Content Platform with a single object storage infrastructure. This supports a wide range of data types, applications, and users with different service level needs in enterprise and cloud environments.

## Software Elements

Table 2 describes the software used to deploy this solution.

TABLE 2. SOFTWARE ELEMENTS

Software	Version
Hitachi Storage Navigator Modular 2	Microcode dependent
SAP HANA	SAP HANA 2.0 SPS01 or later
SUSE Linux Enterprise Server for SAP Applications	12 SP2 or later
VERITAS NetBackup	8.0 (for NetBackup master and media sever)
VERITAS NetBackup SAP HANA Agent	8.0 (for HANA Nodes)
PBS ContentLink	3.3.0.8
Microsoft® Windows® 2012 R2	Standard Edition <ul style="list-style-type: none"><li>For Rack optimized server for solutions</li><li>single node for Management server</li></ul>

## Solution Design

The detailed design for this SAP HANA scale-up backup and recovery reference solution using Backint for SAP HANA with NetBackup includes the following:

- SAP HANA and S/4HANA backup and archive Infrastructure
- Hitachi Content Platform Architecture
- Veritas NetBackup (NBU) Architecture
- PBS ContentLink Architecture
- Fibre Channel SAN Architecture
- Network Architecture
- Backup and Recovery Operations
- Archiving Operations

This paper discusses mainly the architecture of the backup/recovery and archive operations in detail. For configuration details, please refer to the documents from specific vendors.

## SAP HANA and S/4HANA Backup and Archival Infrastructure

### *Backup and Recovery*

For validation of backup and recovery to HCP, this solution uses SAP HANA database running on Hitachi UCP for SAP HANA. However, any database that is supported by SAP and Veritas NBU can be used in a customer implementation. shows the high-level solution design of this reference architecture solution. For details about the architecture and design of such a configuration, please refer to the reference architecture guide [Hitachi Unified Compute Platform 6000 for the SAP HANA Platform in a Scale-up Configuration with Intel Xeon E7-88xx v4 Processors](#).

The following are the high-level details of SAP HANA database backup components and configuration overview:

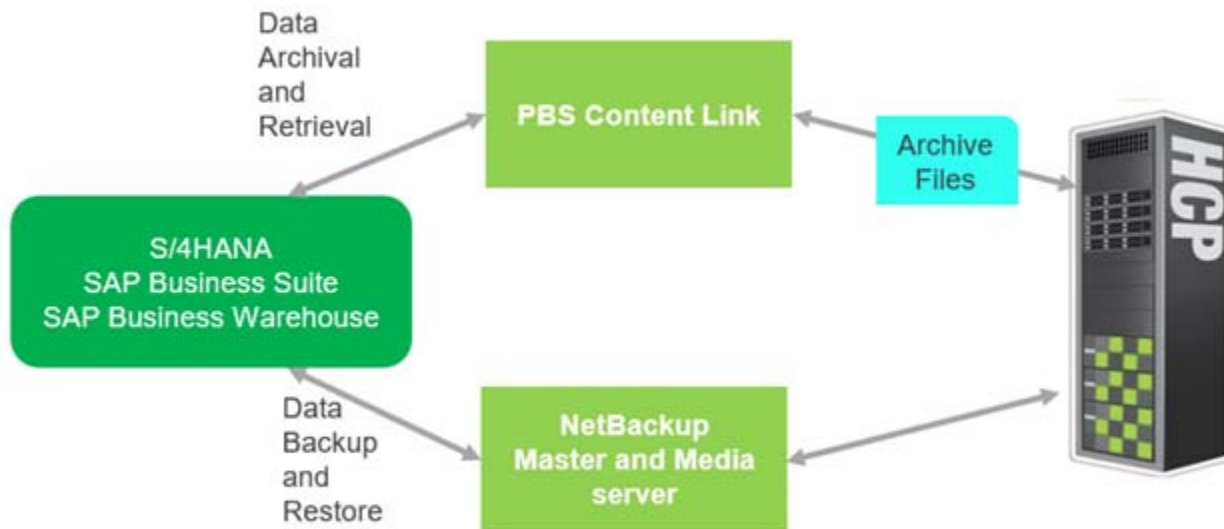
- SAP HANA 2.0 SPS01 is installed on SLES12 SP2 operating system.
- B4H benchmarking data (725 GB) was loaded into the database before performing backup and recovery. This provides a large dataset for measuring backup throughput.
- SAP HANA studio, installed on the management server, was used to perform backup and recovery operations.
- Veritas NBU client for SAP HANA was installed on the SLES12 SP2 server, as described in the Veritas NetBackup (NBU) Architecture section.
- Veritas NBU configuration for SAP HANA can be performed using the administrator guide available at the Veritas website [NetBackup 8.0 for SAP Administrator's Guide](#).
- SAP HANA backup parallelization can be performed for Backint based on the environment as described in [Multistreaming Data Backups with Third-Party Backup Tools](#).
- After the configuration of all the components described in this reference architecture is complete, data backups (full and delta), log backups, and recovery of SAP HANA database to/from HCP can be performed using HANA studio. Backups can be performed on-demand or can be scheduled.

### *Archiving*

For validation of archiving data from SAP applications, this solution uses the S/4HANA application and PBS ContentLink as archiving tool. However, any SAP system that supports archiving using PBS ContentLink can be implemented. The following are the high-level details of S/4HANA archiving components and configuration overview:

- S/4HANA 1610 IDES version was installed for archive testing on SAP HANA 2.0 SPS01.
- S/4HANA and SAP HANA were installed on a single server using the SLES12 SP2 operating system.
- A content repository was defined inside S/4HANA system that connected to HCP using the PBS ContentLink Archive server.
- This content repository can be configured for use by S/4HANA archiving objects to be archived to HCP. The SAP data archiving details and configuration steps can are described at [Introduction to Data Archiving](#).
- After the configuration of all the components described in this reference architecture is complete, SAP data archiving to/from HCP can be performed using standard SAP transactions.

Figure 2



## Hitachi Content Platform (HCP)

In this solution HCP is used as on-premise cloud object storage for storing the backups and archives of the SAP system and SAP HANA database. HCP offers multitenancy which enables users to logically separate the tenants for security and management purposes. Since a single HCP system is used for backup as well as archiving, two separate tenants and namespaces are used for backup and archiving purposes. The following are the high-level details of HCP backup and archiving components and configuration overview:

- Four HCP G10 nodes with 4 TB internal drives and SSD acceleration are used for validation of backup and archive in this solution.
- Total storage capacity in these nodes is about 27 TB, but HCP can scale up to 80 nodes and 800 TB at this time.
- HCP can be installed with internal storage, with separate storage nodes and with external SAN storage. It can be installed on physical servers or virtual machines.
- A separate tenant named "hdbbackup" is created for backup purposes. The namespace for the backup with the name "hdbbucket" is created by Veritas NBU during the configuration.
- A separate tenant "s4hanaarch" and namespace "s4hbucket" is created for archival purposes. The FQDN of the namespace is used during the installation of PBS ContentLink along with other connection details for integration with HCP.

## Veritas NetBackup (NBU) Architecture

Veritas NetBackup is used in this solution as a backup tool. NetBackup is the first backup product to be certified by SAP for protecting SAP HANA. Refer to [Protecting SAP HANA with NetBackup](#) for details about backup and recovery operations of SAP HANA using Backupint provided by NetBackup. The following are the high-level details of NBU components and configuration overview:

- NetBackup master and media server: The two Hitachi Compute Blade 2000 X57A2 server blades with Windows 2012 R2 Standard edition have the VERITAS NetBackup master and media servers installed. The NetBackup master and media server operating system platforms are supported as described in the [NetBackup 8.0 - 8.x.x Operating System Software Compatibility List](#). For detailed steps on how to perform the NetBackup master and media server installation and upgrade refer to the [VERITAS NetBackup 8.0 Installation Guide](#) and [NetBackup 8.0 Upgrade Guide](#).
- NetBackup Client: NetBackup master and media server must be installed before installing the NetBackup client software on the SAP HANA server. As per [Veritas NetBackup™ 8.0 - 8.x.x Database and Application Agent Compatibility List](#), SAP HANA 2.0 SPS01 is supported with SLES12 SP2. For detailed steps on how to perform NetBackup client software install and upgrade on HANA nodes, refer to the [VERITAS NetBackup 8.0 Installation Guide](#) and [NetBackup 8.0 Upgrade Guide](#).
- Using the NetBackup Administrator Console, HCP can be configured as on-premise cloud storage. When configured with appropriate credentials, NBU creates a namespace inside HCP. Use the disk pool created using this connection to create a backup policy. This backup policy is used in the SAP HANA node backup configuration to perform backup/recovery to/from HCP.

During the NetBackup Client software installation, use the SAP HANA backup network name assigned to each of the SAP HANA nodes as the NetBackup client name.

## PBS ContentLink Architecture

PBS ContentLink is used in this solution as an SAP archiving tool. PBS ContentLink (PCL) is a lean software interface that acts as the bridge between SAP and compliant storage systems (on-premise) or virtual cloud storage. PCL covers structured and unstructured data with SAP ArchiveLink and WebDAV. Using PBS ContentLink, customers can archive multiple supported data types from SAP to HCP. Refer to the link [PBS ContentLink - Lean Interface Solution for SAP Data Storage](#) for all supported data types. The following are the high-level details of PBS ContentLink components and configuration overview:

- PBS ContentLink Archive Server: The Hitachi Compute Blade 2000 X57A2 server blade with Windows 2012 R2 Standard edition has PBS ContentLink Archive Server installed.
- PBS ContentLink software integrates with HCP and SAP systems using the credentials supplied during the installation procedure.
- Customers can also configure features like encryption, SSL, and HCP replication during installation.
- Instructions for installation and administration are available within the installation media provided to customers.

## Fibre Channel SAN Architecture

Refer to reference architecture [SAP HANA Tailored Data Center Integration on Hitachi Virtual Storage Platform Gx00 using Hitachi Dynamic Provisioning Pools](#) for details about running multiple SAP HANA scale-up nodes on a single VSP G600. This is a Tailored Data Center Integration (TDI) implementation, so storage connections, design, and layout may change depending on customer requirements.

## Network Architecture

The network architecture in this solution can be divided into three sections:

### *Management Network*

Used for management of the chassis, server blades, VSP G400, and Brocade VDX 6740 switches. This provides a 1 Gb/sec network for management of these hardware components through the Rack optimized server for solutions single node management server.

- This resides on a 1 GbE Brocade ICX 6430-48 port switch.
- The management network does not need to have a VLAN assigned to it.
- The Brocade ICX 6430-48 port switch uses the default switch configuration.
- The Hitachi Compute Blade 2500 and Hitachi Compute Blade 2000 chassis have two management modules to secure fault tolerance. All management modules are connected to the Brocade ICX 6430-48 port switch.
- Each BMC port of HCP G10 node is connected to Brocade ICX 6430-48 switch for management.
- VSP G400 management server (MiniMapp server) is connected to the Brocade ICX 6430-48 port switch.
- The four Brocade VDX 6740 switches are connected to the Brocade ICX 6430-48 port switch.
- The Rack optimized server for solutions single node is connected to the Brocade ICX 6430-48 port switch.

### *Back-end Network*

The isolated back-end network connects the HCP nodes to each other through two redundant back-end Ethernet switches. There are two 10 GbE Intel SFP+ cards on the server that provide redundant 10 Gb Ethernet back-end network connections. Two Brocade VDX 6740 switches are used for this purpose.

Each node has a pair of bonded Ethernet ports that are used to connect the node to these switches. All of the nodes in an HCP system must be on the same system-specific private back-end subnet. This subnet is used only for inter-node communication. It must not be connected to your corporate network, which is used for other types of data communication. This can be configured during the installation of HCP.

Figure 3 shows the information required for back-end network connectivity for the HCP G10 nodes.



Figure 3



- **Red Frames** — Indicates front-end port on 10 GbE Intel SFP+ cards of HCP G10 node
- **Blue Frames** — Indicates back-end port on 10 GbE Intel SFP+ cards of HCP G10 node

Brocade VDX6740 Switch B					
10101	P01	P03	P05	P07	
○ SER	ISL	SP	N01	N03	
○ MG	ISL	R	N02	N04	
□ MT			SEC	SEC	
	P02	P04	P06	P08	

Brocade VDX6740 Switch A					
10101	P01	P03	P05	P07	
○ SER	ISL	SP	N01	N03	
○ MG	ISL	R	N02	N04	
□ MT			PRI	PRI	
	P02	P04	P06	P08	

\*\* Note: Figure only shows first 8 ports of a 48-port Brocade VDX 6740 switch

- SER — the serial port used to configure the switch.
- MGMT — the management port used to configure the switch.
- ISL — the inter-switch link used to connect the two back-end switches to one another.
- SP — the service port used by HDS service personnel.
- R — the port is reserved.
- Nxx PRI - this port is connected to primary port of HCP G10
- Nxx SEC - this port is connected to secondary port of HCP G10

## Front-end Network

The front-end network connects all the components in this solution via a 10 GbE network. Two Brocade VDX 6740 switches are used for this purpose. There are 10 GbE 2-port LAN adapters installed on the PCIe slots of the CB 2500 chassis, CB 2000 chassis, and HCP G10 servers.

- Connect one port of these adapters to each of the Brocade VDX 6740 switch. Figure 3 on page 12 shows the primary and secondary ports for HCP G10 nodes.
- Bond the two ports at the operating system level using active-active network bond mode with options "mode=802.3ad miimon=100 xmit\_hash\_policy=layer3+4 updelay=5000 lacp\_rate=fast" (HCP G10 installation creates the network bond during installation)
- Configure an inter-switch link (ISL) between the two Brocade VDX 6740 switches. Configure the external switch ports with a short LACP timeout value to speed up detecting corrupted connections.

## Backup and Recovery Operations

During normal SAP HANA database operation, data is automatically saved from memory to disk at regular savepoints. Additionally, all data changes are recorded in the redo log. The redo log is saved from memory to disk with each committed database transaction. After a power failure, the database can be restarted like any disk-based database, and it returns to its last consistent state by replaying the redo log since the last savepoint.

While savepoints and log writing protect the data against certain failures like power/memory module failure etc., savepoints do not help if the persistent storage itself is damaged. To protect against data loss due to disk failures or logical database corruption, backups to a secondary storage media such as storage disks or tape libraries are required.

A complete backup of SAP HANA Database performs a backup of the payload (actual) data in the data area. It is important to perform regular automatic backups of the log volume in order to do a point-in-time recovery of the SAP HANA database to the most recent state in the event of failure. Also, it is recommended to perform backup of the configuration files. With Backint, only complete data backup and automatic log backup are performed. Backup of the configuration files needs to be done manually. With the automatic log backup enabled using Backint, the log backups are performed regularly and automatically.

To check the backup tools that can be used with Backint for HANA to create backups for the data and redo log of the SAP HANA system refer to the [SAP Note 1730932](#). In this reference solution, Backint for HANA integration was tested with VERITAS NetBackup.

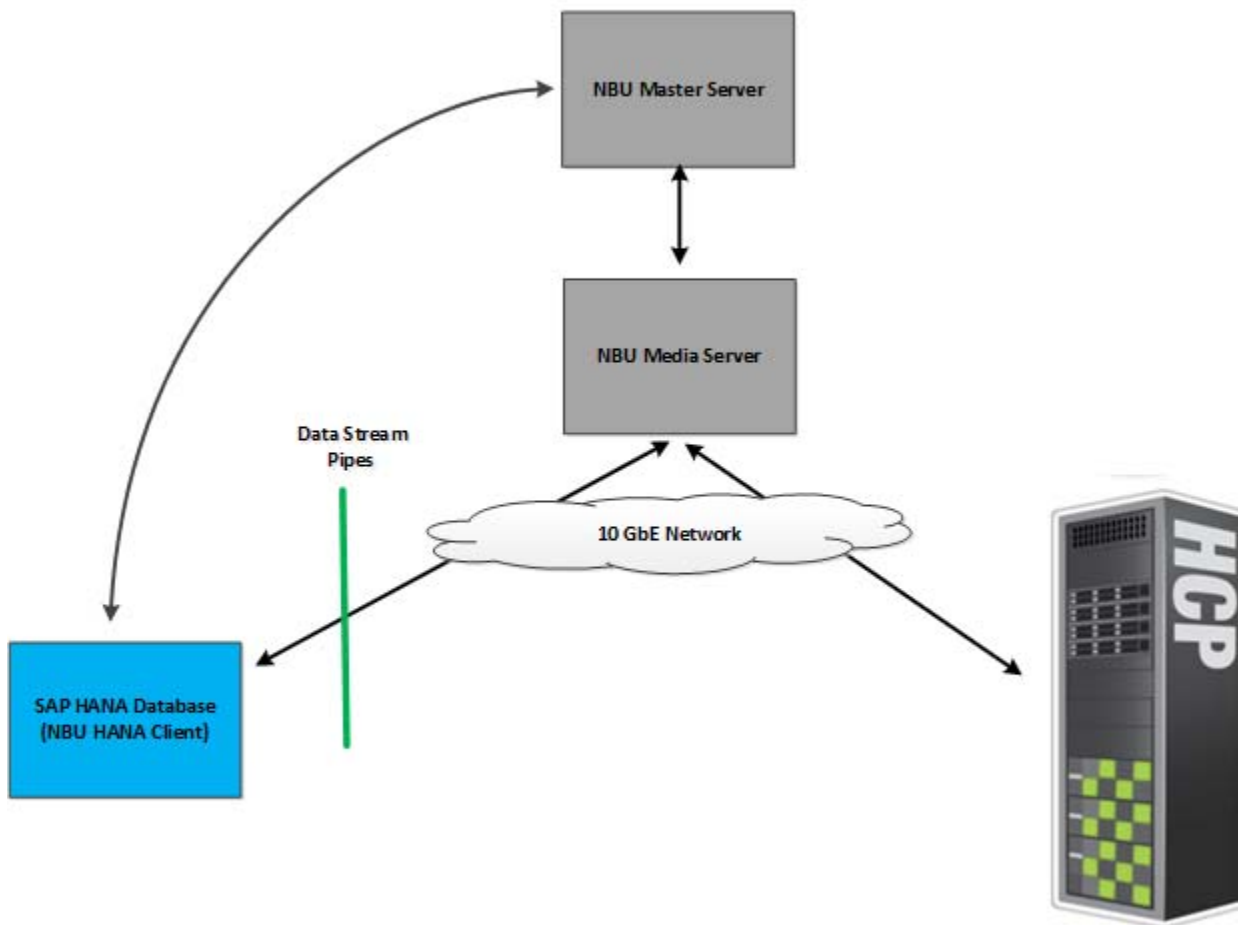
In this solution backup and recovery processes were initiated from the SAP HANA studio. Refer to the [SAP HANA Administration Guide](#) for details on the other scheduling options available to perform the backup and recovery tasks. With Backint, when a backup or recovery task is initiated from HANA studio or an SQL statement, SAP HANA creates a data stream pipe. The SAP HANA agent that runs on the SAP HANA nodes communicates with the NetBackup servers through the Backint interface using these data stream pipes to perform backup or recovery processes.

During backup, each of the SAP HANA services has its own stream pipe for writing backup data. NetBackup SAP HANA agent reads the data stream from these pipes and passes them to the NetBackup server, which is saved to secondary storage (**Hitachi Content Platform**) based on the policy storage chosen as the destination in the NetBackup server backup policy. Finally, SAP HANA transmits backup catalog information, before the NetBackup HANA agent writes a file reporting the result and administrative information such as backup identifiers.

When a recovery task is initiated from SAP HANA studio or an SQL statement, SAP HANA establishes communication with the NetBackup HANA agent using stream pipes and requests the backup data from the NetBackup server. NetBackup HANA agent then streams the backup data received from the NetBackup server through the stream pipes to the SAP HANA services for performing the recovery. As a final step, the NetBackup SAP HANA agent writes a file reporting the results of the operation. Refer to Figure 4 for the backup and restore operations.



Figure 4



## Archive Operations

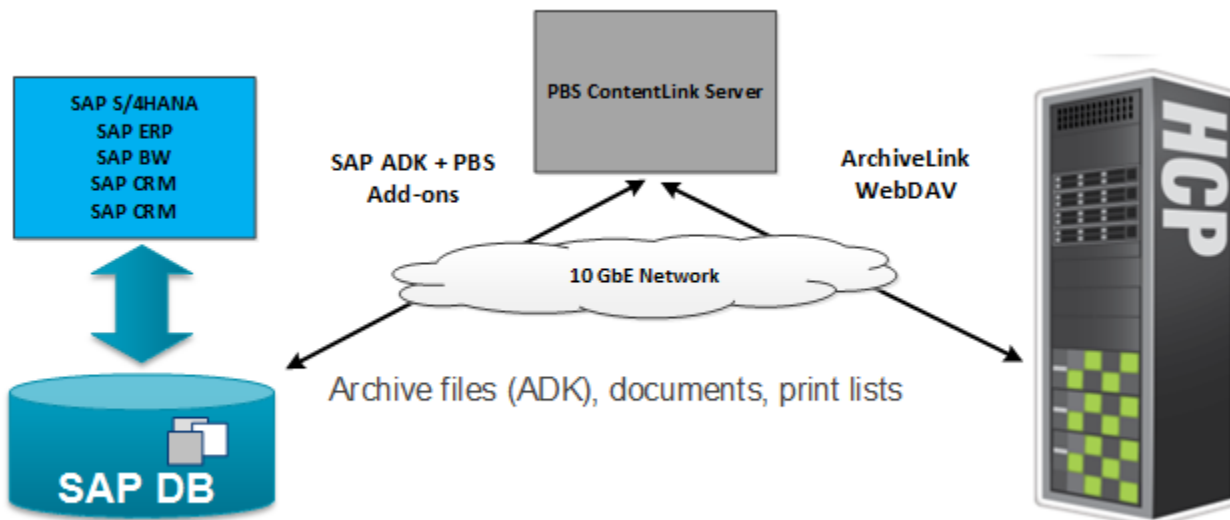
With SAP Archive Solution (PBS Content Link) you can create a direct and lean connection between SAP systems and Hitachi HCP, and in this way enable archiving that meets legal requirements and a future-proof Information Lifecycle Management concept. PBS software is installed on an intermediary server that sits between SAP and HCP. The PBS server has a WebDAV interface to terminate ILM 3.0 traffic, and redirect it to HCP using the REST protocol. The PBS software doesn't save any information; it's stateless, and commits writes only when HCP commits them. PBS Content Link also simultaneously supports the new SAP-specific WebDAV 3.0 protocol in addition to the already well-established SAP ArchiveLink interface. PBS Content Link has been certified by SAP for the integration scenario BC-ILM 3.0 (Information Lifecycle Management for SAP). The WebDAV protocol ensures the usage under the SAP ILM components Retention Manager and Retention Warehouse. PBS Content Link software (Archive Server) runs as a service on major platforms (Windows, Linux, AIX, Solaris) using the platform specific webserver as the underlying engine (IIS for Windows, Apache httpd on the \*NIX platforms).

The PBS ContentLink archive server can be integrated with the SAP system and HCP during the installation. Once the configuration in the SAP system is also complete, archiving of all supported data types can be performed. Below are the steps listed for only for archiving SAP data to HCP. All the steps can be performed through a SAP Archive Administration (SARA) transaction.

- **Create Archive Files:** This step writes the archive files from the SAP system to HCP. To create archive files, you can select the archive object, create variants to define the selection criteria for data to be archived, and schedule the job.
- **Deleting Archived Data from the Database:** The delete program can run automatically or manually after the archiving is finished. It deletes the archived data from the SAP database.
- **Reading Archive Files:** In addition to the data in the online database, archived data can also be read or analyzed. When data is analyzed, it is read sequentially and elementary information such as item number, customer, and order date is displayed in list form. It is possible to run an analysis for one or several archiving sessions. Not all archiving objects offer an analysis program.
- **Reloading Data:** The reloading function is only meant for emergency cases, in which the wrong data was archived due to errors during document selection or customizing. Therefore, data should only be reloaded immediately after the erroneous archiving session. If you reload the data later you may end up with inconsistencies in your database.

The complete archiving procedure is described in [Archiving Procedure at the SAP website](#). Figure 5 shows the archiving operations overview.

**Figure 5**



## Engineering Validation

Validation of the SAP HANA scale-up backup and recovery using Backint for SAP HANA with NetBackup reference solution was conducted in the Hitachi Vantara Lab.

### Test Methodology

#### *Backup Test Methodology*

The validation testing includes the backup and recovery operations discussed in this reference architecture paper. B4H benchmarking data (725 GB) was loaded into the database before performing backup and recovery. This provided a large dataset for measuring backup throughput.

Successful Complete, Differential, and Incremental backups were performed along with log backups to the Hitachi Content Platform. Similarly, successful recovery of the SAP HANA database from backup data stored on Hitachi Content Platform were performed. In addition, the recovery option with all the recovery types (like Recover the database to its most recent state, Recover the database to the following point in time, Recover the database to a specific data backup) were performed from HANA studio.

Backup throughput of 1.1 TB/hour was achieved with this dataset and four HCP G10 nodes. Please note that this is not the maximum ingestion speed for HCP. Throughput can be increased further by taking multiple backups from different SAP HANA databases at the same time and scaling HCP.

#### *Archive Test Methodology*

The validation of archiving includes end to end archive and reload of IDOC archiving objects from S/4HANA SAP system. S/4HANA 1610 IDES version was installed in the Hitachi Vantara lab for this purpose. A test variant was created to archive all the IDOCs created in the year 2016.

- Write (Writes the archive file to HCP)
- Delete (Delete the EDID\* table entries from S/4HANA database)
- Read (Read the archive file created in above step)
- Reload (Reloads the data from HCP back to S/4HANA database to EDIDC\* tables)

## For More Information

Hitachi Vantara Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the [Services](#) website.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the [Resources](#) website.

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantara [Training and Certification](#) website.

For more information about Hitachi products and services, contact your sales representative, partner, or visit the [Hitachi Vantara](#) website.

## Hitachi Vantara



Corporate Headquarters  
2845 Lafayette Street  
Santa Clara, CA 96050-2639 USA  
[www.HitachiVantara.com](http://www.HitachiVantara.com) | [community.HitachiVantara.com](http://community.HitachiVantara.com)

Regional Contact Information  
**Americas:** +1 408 970 1000 or [info@hitachivantara.com](mailto:info@hitachivantara.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hitachivantara.com](mailto:info.emea@hitachivantara.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hitachivantara.com](mailto:hds.marketing.apac@hitachivantara.com)

© Hitachi Vantara Corporation 2017. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation. All other trademarks, service marks and company names are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Vantara.

SL-008-00, 2017.