HITACHI
Inspire the Next

TECHNICAL PAPER

Hitachi Unified Compute Platform for the SAP HANA Platform in a Scale-Up Configuration Using Hitachi Compute Blade 2500: Global File System 2 Active-Standby High Availability Solution

Reference Architecture Guide

By David Pascuzzi

October 2015

Hitachi Data Systems

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.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.

Contents

Hitachi Unified Compute Platform for the SAP HANA Platform in a Scale-Up Configuration Using Hitachi Compute Blade 2500: Global File System 2 Active-Standby High Availability Solution

Reference Architecture Guide

This reference architecture guide describes how to extend Hitachi Unified Compute Platform (UCP) for the SAP HANA Platform in a scale-up configuration using Hitachi Compute Blade 2500 (CB 2500) for different sized solutions to provide high availability (HA) using Global File System 2 (GFS2) and SAP HANA Storage API with the following additional items:

- 520X B2 server blades (for CB 2500 chassis)
- SMP connection board for 520X server blades
- LAN PCIe adapters
- Fibre Channel adapters
- RedHat Resilient Storage Add-On
- License for RedHat Resilient Storage Add-On for each blade

Unified Compute Platform for SAP HANA is a preconfigured analytical appliance, ready to plug into a network to provide real-time access to operational data for use in analytic models.

This standby node extension to the appliance supports the configurations listed in Table 1. The configuration for the standby node must match the configuration used for the SAP HANA node.

Table 1. Supported Scale-Up Configuration Sizes

Description	2-Socket	4-Socket	8-Socket
Number of CPUs	2	4	8
Number of Server Blades	1	2	4
Memory	■ 128 GB ■ 256 GB ■ 384 GB ■ 512 GB ■ 768 GB ■ 1024 GB ■ 1536 GB	■ 128 GB ■ 256 GB ■ 384 GB ■ 512 GB ■ 768 GB ■ 1024 GB ■ 1536 GB ■ 2048 GB ■ 3072 GB	■ 256 GB ■ 512 GB ■ 768 GB ■ 1024 GB ■ 1536 GB ■ 2048 GB ■ 3072 GB ■ 4096 GB ■ 6144 GB

This technical paper assumes that you have familiarity with the following:

- Storage area network (SAN)-based storage systems
- General storage concepts
- Common IT storage practices
- SAP HANA
- RHEL system administration
- RHEL Resilient Storage cluster administration
- Hitachi UCP for SAP HANA configuration

Solution Overview

This reference solution extends an existing appliance by adding the following:

- 520X B2 server blades for Hitachi Compute Blade 2500 chassis
- SMP connection board for 520X server blades
- LAN PCIe adapters
- Fibre Channel adapters

Figure 1shows the changes to the existing topology for this reference architecture.

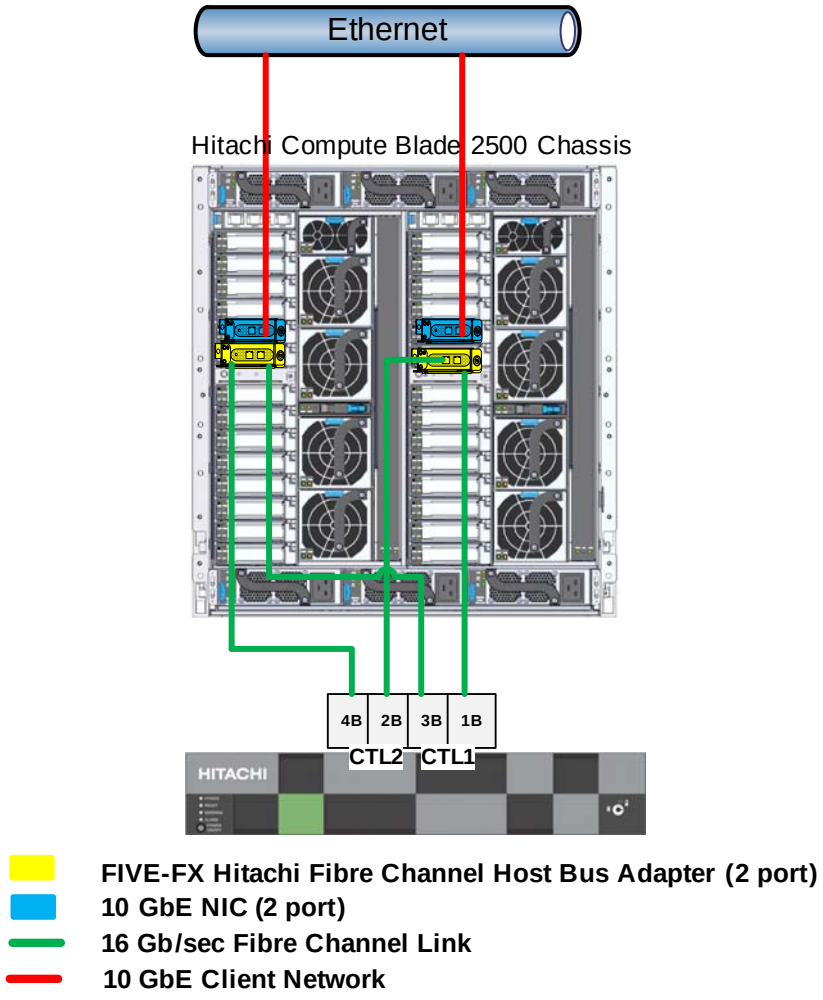

Figure 1

Key Solution Elements

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

Hardware Elements

Table 2 lists the hardware used to deploy the specific scale-up configuration of Unified Compute Platform for SAP HANA for the different sized solutions.

Table 2. Hardware Elements

Hardware		Quantity	Configuration (per unit)	Role
520X B2 server blade	2-Socket	1	■ 2 × 18-core processors ■ RAM for all memory sizes listed in Table 1 ■ 1 × 2-port pass-through mezzanine card on mezzanine slot 2 and mezzanine slot 4 of blade 9 and blade 10 (on single blade for 2-Socket)	SAP HANA server
	4-Socket	2		
	8-Socket	4		
SMP connection board for 520X server blade	2-Socket	0	■ SMP connection board ■ 2-blade SMP connection board for 4-Socket ■ 4-blade SMP connection board for 8-Socket ■ SMP expansion module ■ SMP connector cover	SMP connector
	4-Socket	1		
	8-Socket	1		

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.

The configuration uses one, two, or four 520X B2 server blades in the existing Hitachi Compute Blade 2500 chassis for the different sized solutions listed in Table 1 on page 1.

The configuration for the 520X B2 server blades chosen for the standby node must match the configuration for the active SAP HANA node.

Table 3 lists the configurations for the 520X B2 server blades used in the various sized solutions.

Table 3. 520X B2 Server Blade Configuration

Feature		Configuration		
		Small (2-Socket)	Medium (4-Socket)	Large (8-Socket)
Processors		2 Intel Xeon E7-8880 processors	4 Intel Xeon E7-8880 processors	8 Intel Xeon E7-8880 processors
Processor SKU		Intel Xeon E7-8880 v3		
Processor frequency		2.3 GHz		
Processor cores		18 cores		
Number of blades		1	2	4
Number of DIMMs per blade	128 GB	8 × 16 GB DIMMs	8 × 8 GB DIMMs	N/A
	256 GB	16 × 16 GB DIMMs	8 × 8 GB DIMMs	8 x 8 GB DIMMs
	384 GB	24 × 16 GB DIMMs	24 × 8 GB DIMMs	N/A
	512 GB	32 × 16 GB DIMMs	16 × 16 GB DIMMs	16 × 8 GB DIMMs
	768 GB	48 × 16 GB DIMMs	24 × 16 GB DIMMs	24 × 8 GB DIMMs
	1024 GB	32 × 32 GB DIMMs	32 × 16 GB DIMMs	16 × 16 GB DIMMs
	1536 GB	48 × 32 GB DIMMs	48 × 16 GB DIMMs	24 × 16 GB DIMMs
	2048 GB	N/A	32 × 32 GB DIMMs	32 × 16 GB DIMMs
	3072 GB	N/A	48 × 32 GB DIMMs	48 × 16 GB DIMMs
	4096 GB	N/A	N/A	32 × 32 GB DIMMs
	6144 GB	N/A	N/A	48 × 32 GB DIMMs
Network Ports		2 × 2-port 10GBASE-SR LAN PCIe adapter on two I/O board modules: ▪ IOBD 09B ▪ IOBD 10B		
Fibre Channel Ports		2 × Hitachi 16 Gb/sec 2-port Fibre Channel adapters on two I/O board modules: ▪ IOBD 09A ▪ IOBD 10A		

Figure 2 shows the layout of the I/O board modules from the back of the Hitachi Compute Blade 2500 chassis.

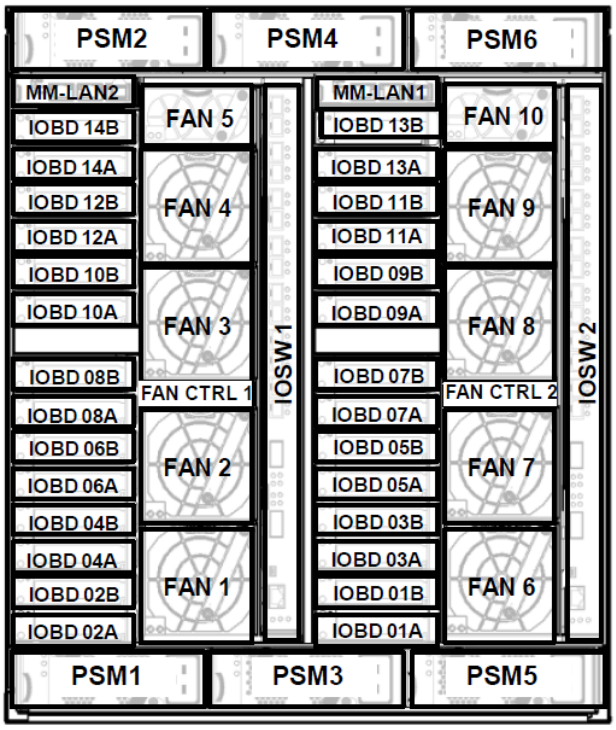

Figure 2

Symmetric Multiprocessing Connector

For multiple server blades, the solution uses symmetric multiprocessing (SMP) technology to combine multiple server blade resources into a single server.

The 520X B2 server blades use SMP in one of two ways, depending on the size of the solution:

- Combine two server blades with a 2-blade SMP connection board for the 520X blade server
- Combine four server blades with a 4-blade SMP connection board for the 520X blade server

Hitachi Virtual Storage Platform

This solution uses the existing Hitachi Virtual Storage Platform (VSP) from the base solution. The operating system LUNs, SAP HANA data, and log LUNs reside on this storage array, as well as the shared LUNs for SAP HANA configuration files, binaries and traces. This solution creates an extra LDEV on the parity group used for the operating system.

Software Elements

Table 4 describes the software products used to deploy the standby node.

Table 4. Software Elements

Software	Version
RedHat Enterprise Linux for SAP	RHEL 6.6 or later
RedHat Resilient Storage Add-on	RHEL 6.6 or later
SAP HANA Platform	1.0 SPS9, Rev. 97 or later

SAP HANA

The SAP HANA platform is flexible, multipurpose in-memory software. It combines SAP software components optimized to specific hardware. These components come from leading hardware partners of SAP, including Hitachi Data Systems.

The SAP HANA platform allows customers to analyze large volumes of data in real time. It is also a development platform, providing an infrastructure and tools for building high-performance applications based on SAP HANA Extended Application Services (SAP HANA XS). It is the foundation of various SAP HANA editions, like the SAP HANA Platform Edition, providing core database technology, and the SAP HANA Enterprise Edition, bundling additional components for data provisioning. The SAP HANA Platform Edition integrates a number of SAP components, including the SAP HANA database, SAP HANA studio and SAP HANA clients.

As a SAP customer, you can get more information on the SAP HANA platform at the SAP Service Marketplace and help.sap.com.

See **Installation and Upgrade Information** on SAP HANA Platform for the following guides and all other SAP-related documentation:

- **The SAP HANA Master Guide** — This guide is the entry point for planning the installation of your SAP HANA system landscape
- **The SAP HANA Master Update Guide** — This guide describes how to update the SAP HANA platform and related components
- **SAP HANA Server Installation and Update Guide** — This guide describes how to install and update an SAP HANA system with the SAP HANA life-cycle management tools
  - There are various related installation guides to install the required SAP in-memory database and the other software components for the different replication technologies
- **SAP HANA Technical Operations Manual** — This guide is the entry point for administering and operating your SAP HANA system landscape

SAP Integration and Certification Center (SAP ICC) provides information about SAP-certified SAP HANA appliances by SAP hardware partners.

Red Hat Enterprise Linux (RHEL)

Hitachi Unified Compute Platform for the SAP HANA Platform in a scale-up configuration can run on the following:

- 64-bit Red Hat Enterprise Linux for SAP with Resilient Storage Add-on

The SAP HANA standby node must run the same operating system version and SAP HANA Version as the SAP HANA worker node.

The initially delivered configuration of the operating system should persist. Changing the configuration settings can cause significant performance problems to occur.

Do not make any modifications to the operating system, except as noted or approved by SAP. To modify your operating system configuration or the installed software packages, follow the guidelines given by SAP and the operating system distributor. See the following SAP Notes and RHEL:

2136965 - SAP HANA DB: Recommended OS settings for RHEL 6.6

For more details, see the section “Updating and Patching the Operating System,” in the SAP HANA Technical Operations Manual.

Solution Design

The detailed design for this extension to Hitachi Unified Compute Platform for the SAP HANA Platform in a scale-up configuration includes the following:

- Hitachi Compute Blade 2500 Chassis Configuration
- 520X B2 Server Blade Architecture
- Fibre Channel Architecture
- Network Architecture
- Storage Architecture
- Parity Group, RAID, and LUN Configuration
- SAP HANA Configuration

Hitachi Compute Blade 2500 Chassis Configuration

Figure 3 shows the front and back view of the Hitachi Compute Blade 2500 chassis for different configuration sizes.

Small (2-Socket) – One Blade

Hitachi Compute Blade 2500 Chassis 1 (Front)

Medium (4-Socket) – Two Blades (SMP)

Hitachi Compute Blade 2500 Chassis 1 (Front)

Large (8-Socket) – Four Blades (SMP)

Hitachi Compute Blade 2500 Chassis 1 (Front)

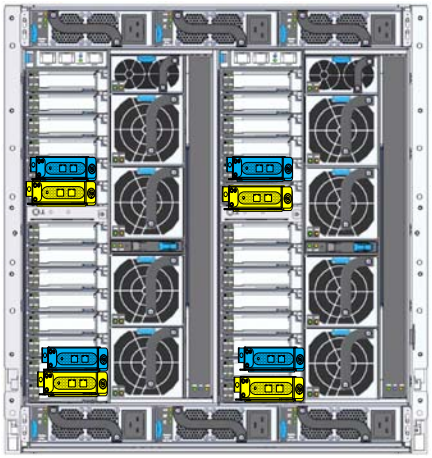

Hitachi Compute Blade 2500 Chassis 1 (Back)

FIVE-FX Hitachi Fibre Channel Host Bus Adapter (2 port)

10 GbE NIC (2 port)

Figure 3

The solution uses the existing Hitachi Compute Blade 2500 chassis with the following components added:

- One to four 520X B2 server blades. Refer to Table 1 on page 1
- A maximum of 28 I/O board modules (IOBD) can be mounted on one Hitachi Compute Blade 2500 chassis; the solution extends the configuration by four I/O board modules.
  - Hitachi 16 Gb/sec 2-port Fibre Channel PCI-X adapters are installed on IOBD 09A and IOBD 10A
  - 10GBase-SR 2-port network PCI-Ex adapters are installed on IOBD 09B and IOBD 10B

520X B2 Server Blade Architecture

Each solution size has a different number of server blades, from one up to four full-width server blades. Multiple server blades use a two- or four-server blade SMP connection board to create a single four- or eight-socket SMP node with a total of 72 or 144 cores and different memory sizes.

Table 5 lists the server blade configuration details for small, medium, and large solution sizes.

Table 5. Server Blade Configuration

	Small (2-Socket)	Medium (4-Socket)	Large (8-Socket)
Server Blades	Total of 1 server blade: - Blade 9 (primary)	Total of 2 server blades: - Blade 9 (non-primary) - Blade 11 (primary)	Total of 4 server blades (bottom-up): - Blade 9 (non-primary) - Blade 11 (non-primary) - Blade 13 (non-primary) - Blade 14 (primary)
Total Number of CPU Cores	36	72	144

Fibre Channel Architecture

As shown in Table 6, the solution uses two Hitachi 16 Gb/sec 2-port Fibre Channel Adapters installed on the PCIe slot of the I/O board module of blade 9. This solution uses four 16 Gb/sec Fibre Channel ports on the Virtual Storage Platform directly attached to the Hitachi Compute Blade 2500 server chassis via the Fibre Channel PCIe adapters.

Table 6. Fibre Channel Port Mapping

Blade #	PCIe Slot #	Port #	VSP Port
Blade 9	IOBD 09A	0	1B
		1	2B
	IOBD 10A	0	3B
		1	4B

This configuration supports high availability by providing multiple paths from the host within Hitachi Compute Blade 2500 to multiple ports on Virtual Storage Platform.

Figure 1 shows the direct connect Fibre Channel architecture for this solution.

Set the port properties for the point-to-point connection between Hitachi Compute Blade 2500 and Virtual Storage Platform, as shown in Table 7.

Table 7. Port Properties

Property	Value
Port Security	**Disabled**
Port Speed	**Auto (16 Gb/sec)**
Fabric	**ON**
Connection Type	**P-to-P**

Network Architecture

There are two 10GBASE-SR 2-port LAN adapters installed on the PCIe slots of the I/O board module of blade 1 of the Hitachi Compute Blade 2500 chassis. The solution uses two 10 GbE ports on the 10GBASE-SR 2-port LAN adapters for connectivity with the 10 GbE external switches.

Make the following network connections for the client network setup of the SAP HANA node as an uplink network setup:

- Connect the following to the external switches:
  - Port 0 of the I/O board module on PCIe slot IOBD 09B
  - Port 0 of the I/O board module on PCIe slot IOBD 10B
- Bond the corresponding two ports eth9901 and eth9902 as bond0 at the operating system level using the `active-active network bond mode` command with options "mode= 802.3ad miimon=100 xmit_hash_policy=layer3+4 updelay=5000 lacp_rate=fast", which acts as the client network for the SAP HANA node
- Configure the external switch ports with a short LACP timeout value to speed up detecting corrupted connections.

The compute network setup uses the ports on the 10GBASE-SR 2-port LAN adapters, as listed in Table 8.

Table 8. Network Setup Using 10GBASE-SR 2-Port LAN Adapter

Server Blade	PCIe Slot Number	Switch Module Port	Network Description
Blade 9	IOBD 09B	0	Client network for SAP HANA node
		1	Free for use as uplink network
	IOBD 10B	0	Client network for SAP HANA node
		1	Free for use as uplink network

Figure 1 shows the standard network configuration used for Hitachi Compute Blade 2500 chassis for this solution.

Storage Architecture

Hitachi Virtual Storage Platform is used for this SAP HANA solution. Sizing and configuring of storage including storage drive box trays (DBS), spare drives, the operating system volume (OS), SAP HANA shared volume (/hana/shared), SAP HANA log volume (/hana/log), and SAP HANA data volume (/hana/data), varies for different sizes by taking into account requirements such as I/O and capacity.

For this solution, the changes to the existing storage are the addition of a 100 GB LUN on the operating system (OS) RAID and a 50 MB LUN on the OS RAID. For shared, data, and log storage this solution uses the existing storage.

Parity Group, RAID, and LUN Configuration

This reference architecture uses various configurations of parity group, RAID level, and LUNs on Hitachi Virtual Storage Platform for different sized solutions.

Table 9 shows the parity groups and LDEV assignments that are applied to all sized memories listed in Table 1 on page 1. These LDEV are added to the existing OS parity group.

Table 9. Parity Group and LDEV

Parity Group	RAID Level and Disks	LDEV ID	LDEV Size	HOST ID	MPU ID	Description	Size Applied
1	RAID-6 (6D+2P) on 600 GB 10k RPM SAS drives	0:01:05	100 GB	000	MPU10	Operating system boot volume	All
		0:01:06	50 MB	Next available, after existing shared data and log volumes.	MPU11	STONITH device	All

Besides using the new OS and STONITH LDEVs, this solution uses the existing shared, data, and log volumes. When the stand by node becomes active, the SAP HANA Storage API will mount the storage to it.

The STONITH device is added to the parity groups for both nodes. It should be assigned the next available host ID.

SAP HANA Configuration

This section explains the SAP HANA configuration.

Operating System Boot Configuration

The SAP HANA scale-up configuration requires a boot volume. It uses one 100 GB LUN on Hitachi Virtual Storage Platform for the operating system volume for the SAP HANA appliance.

The four ports of Hitachi 16 Gb Fibre Channel 2-port adapters have the 100 GB operating system LUN configured as the primary boot device. The operating system LUN holds partitions for RedHat Enterprise Linux on the OS LUN */usr/sap/* and the Linux swap space.

Activate Device-Mapper Multipath

This reference architecture uses Device-mapper Multipath, a component of the native Linux operating system.

Using Device-mapper Multipath allows the configuration of multiple I/O paths between the server blades, and Hitachi Virtual Storage Platform. Multipathing aggregates all physical I/O paths into a single logical path. The LUNs are always available, unless all four paths fail.

Device-mapper Multipath is used for the following I/O paths:

- SAP HANA server operating system LUN
- SAP HANA data volume LUN
- SAP HANA log volume LUN
- SAP HANA shared volume LUN
- RHEL Resilient Storage STONITH LUN

GFS2

This reference architecture uses Global File System (GFS2), Pacemaker, and Clustered Logical Volume manager (cLVM). These are components of the RHEL with Resilient Storage add-on operating system. They provide a shared clustered file system across nodes without using an NFS server.

GFS2 has the following limitations:

- Maximum tested file system size is 100 TB
- Maximum number of nodes is 16

The SAP HANA shared volume (/hana/shared) will be converted to a GFS2 shared file system.

STONITH Configuration

This solution requires a STONITH device to control access to clustered resources. The LUN used will not be formatted or mounted. It needs to be accessible from both nodes. Unlike the rest of the storage used in this solution, the STONITH device will not have a file system created on it.

SAP HANA Data Volume Configuration

This solution uses the data volume from the SAP HANA worker node. The standby node will need the data volume's LUNs added to its ports.

SAP HANA Log Volumes Configuration

This solution uses the log volume from the SAP HANA worker node. The standby node will need the log volume's LUNs added to its ports

SAP HANA Shared Volume

This solution uses the shared volume from the SAP HANA worker node. The standby node will need the shared volume's LUNs added to its ports. This volume will be modified to support the GFS2 file system.

SAP HANA Appliance Software Installation

This solution uses the SAP HANA software components that were installed in the shared directory during the SAP HANA worker node installation. The only component that will be installed on the standby node is the SAP Host agent.

This solution uses the SAP landscape management tools on the SAP HANA worker node to install the standby node. The standby node will access the SAP HANA database and SAP HANA client software from the shared volume.

For More Information

Hitachi Data Systems 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 Hitachi Data Systems Global Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Education website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems

Corporate Headquarters
2845 Lafayette Street
Santa Clara, CA 96050-2639 USA
www.HDS.com community.HDS.com

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

AS-433-00 October 2015.