



**Protecting Oracle® Data Integrity:  
Hitachi Database Validator  
Technology and the Hitachi  
Freedom Storage™ Lightning  
9900™ V Series Systems**

A Technical White Paper

by David Broom and Rasha Hasaneen

January 2003

---

## Executive Summary

The cooperative partnership between Hitachi, Ltd./Hitachi Data Systems, the leading storage provider, and Oracle Corporation, the leading database vendor, has delivered one of the most robust, and fully protected validation database platform solutions available today. The new data protection technology was designed specifically for Oracle8i™ or Oracle9i™ databases to prevent data corruption between the database and the networked storage disk system. The technology, implemented at the microchip and microcode level, is integrated with high-end Hitachi Freedom Storage™ Lightning 9900™ V Series systems.

The new data protection technology was initiated as a result of the Hitachi Data Systems decision to support the Oracle® Hardware Assisted Resilient Data (H.A.R.D.) Initiative. A global initiative that constructs resilient data management and storage solutions, H.A.R.D was launched by Oracle in November 2001. Engineering development groups from Hitachi Data Systems and Oracle worked together to integrate Oracle's data integrity checksum algorithms into the Lightning 9900 V Series systems. Together, both companies have created technology that prevents corrupted data from being written into Oracle's internal file types: Oracle database files, REDO log files, and control files.

Joint engineering resulted in the Hitachi Database (DB) Validator product, one of the optional functions in Lightning 9900 V Series systems, which prevents corrupted data blocks generated in the database-to-storage system infrastructure from being written onto the storage disk. Data corruption can occur in any or all layers between the database management system and the storage system before the data is written into the storage system, i.e. while passing through the operating system, channel adapter, or network. While this type of data corruption rarely occurs, once corrupted data is written into the storage system, it is difficult and time-consuming to detect the underlying cause, restore the system, and recover the database. DB Validator technology helps prevent potentially disastrous data corruption and minimizes the risk and potential cost of backup, restore, and recovery operations. Lightning 9900 V Series systems and Oracle8i or Oracle9i databases provide a resilient solution that enables the storage system to operate 24/7—the uptime required by today's enterprises.

In May 2000, DB Validator technology testing was conducted at the Hitachi-Oracle SAN Solution Technology Center (SSTC), established at the Oracle offices in Japan. At the SSTC, special analysis focused on performing collaborative verification testing of actual customer operations and showcasing integration and validation of advanced Hitachi Freedom Storage system capabilities with the Oracle database. Technical data verified at the SSTC will be translated and provided worldwide, which will enable Hitachi/Hitachi Data Systems and Oracle to provide better and more tightly integrated solutions to the worldwide market.

The data validation functionality Hitachi has developed with Oracle will further ensure data integrity and provide more secure storage systems to our joint customers. Hitachi/Hitachi Data Systems plan to continue to working closely with Oracle to deliver additional service-ready solutions in the future.

The purpose of this technical white paper is to profile the collaboration of Hitachi, Ltd./Hitachi Data Systems and Oracle, the resulting DB Validator technology, and its data corruption prevention capabilities.

---

## Contents

<b>Introduction</b>	<b>1</b>
<b>Oracle's H.A.R.D. Initiative</b>	<b>2</b>
<b>DB Validator (Implementation of Oracle H.A.R.D.) by Hitachi Data Systems</b>	<b>2</b>
<b>The Collaborative Relationship</b>	<b>3</b>
<b>Business Benefits of DB Validator and Hitachi TrueNorth™ Vision and Strategy</b>	<b>3</b>
Hitachi TrueNorth™	3
Hitachi Freedom Storage Software	3
Simplify, Protect, and Optimize	3
<b>DB Validator Technology Details</b>	<b>4</b>
<b>DB Validator Solution Details</b>	<b>4</b>
<b>DB Validator Service Details</b>	<b>5</b>
Professional Services	5
Customer Service and Support	5
<b>DB Validator Functionality</b>	<b>5</b>
Database System and Data Corruption	5
Validation Mechanism	6
<b>Summary</b>	<b>8</b>

# Protecting Oracle® Data Integrity: Hitachi Database Validator Technology and the Hitachi Freedom Storage™ Lightning 9900™ V Series Systems

A Technical White Paper

by David Broom and Rasha Hasaneen

---

## Introduction

Data flows through many software and hardware layers on its way from the application to the disk storage system. The Oracle® database, at the origin of the I/O data path, and the Hitachi Freedom Storage™ Lightning 9900™ V Series system, at the termination of the I/O data path, have no control over the software and hardware layers in between. Both Oracle databases and Lightning 9900 V Series systems protect data from corruption within their own environments. For example: Lightning 9900 V Series systems already perform extensive checksum-type validation of data integrity from the point of entry through the entire system.

With exhaustive testing, and even after extensive production use, obscure and “rare case” bugs can sometimes occur in intermediate software and hardware layers. Such rare case bugs have, in some cases, caused corruption of user data. This corruption can go undetected for an extended period of time; however, when discovered it can result in significant outage to business-critical applications, and demand time and energy of technical resources to identify and correct the problem. As with most recovery processes, the time to recover is directly correlated with the length of time in between the corruption actually occurring and detection of that corruption. Therefore, business functions that are data-dependent must quickly eliminate these rare bugs.

Together, Hitachi, Ltd., Hitachi Data Systems, and Oracle have created technology that prevents corrupted data from being written into Oracle’s internal file types: Oracle database files, REDO log files, and control files. The combination of Hitachi Database (DB) Validator technology and Oracle Hardware Assisted Resilient Data (H.A.R.D.) Initiative detects corrupt data as it arrives at the storage system, and then rejects the I/O operation. This prevents the application’s Oracle database transaction from appearing to complete successfully, stopping corrupt data from ever being written to the database in the first place. The DB Validator function immediately sends a notification to the appropriate technical engineers for identification and resolution.

---

Without DB Validator technology, corrupt data would only be discovered when the data was read back at a later date—when it may be very difficult or even impossible to reconstruct. And worse yet, the absence of the DB Validator function could result in the data corruption not being detected at all.

## **Oracle's H.A.R.D. Initiative**

The Oracle H.A.R.D Initiative is a comprehensive program designed to prevent data corruption before it actually occurs. While data corruption is rare, it can have catastrophic effects on databases and, ultimately, on businesses. By implementing Oracle data validation algorithms inside the storage devices, Oracle will prevent corrupted data from ever being permanently written to the storage system. This type of end-to-end, high-level software to low-level hardware validation has never been implemented before. H.A.R.D will eliminate a large percentage of failures that the database industry has, up to this point, been powerless to prevent. Redundant Array of Independent Disks (RAID) has gained a wide following in the storage industry by ensuring the physical protection of data. H.A.R.D takes data protection to the next level by going beyond protecting physical bits to protecting business data.

The classifications of data corruption addressed with H.A.R.D include:

- Writes of physically and logically corrupt blocks
- Writes of blocks to incorrect locations
- Erroneous writes by non-Oracle programs to Oracle data
- Partially written blocks
- Writes from cluster nodes that are considered dead
- Lost writes
- Corrupted third-party backups

The H.A.R.D Initiative includes several technologies that can be embedded in storage devices to prevent all these types of corruption. The Hitachi/Hitachi Data Systems implementation of Oracle H.A.R.D (DB Validator technology) addresses all kinds of data corruption.

## **DB Validator (Implementation of Oracle H.A.R.D.) by Hitachi Data Systems**

DB Validator is a function within high-end Hitachi Freedom Storage systems that is specifically designed to prevent data corruption between the Oracle database platform and the storage system. This function prevents corrupted data blocks, generated in the database-to-storage system infrastructure, from being written onto the storage disk. The conventional combination of networked storage and database management software has a risk of data corruption while data is being written on the storage system. This type of data corruption rarely occurs; however, once corrupted data is written into storage, it is difficult and time-consuming to detect the underlying cause, restore the system, and recover the database. DB Validator technology helps prevent potentially disastrous corrupted data environments, and it minimizes risk and the potential costs associated with backup, restore, and recovery operations.

---

DB Validator began with the Hitachi/Hitachi Data Systems decision to support Oracle H.A.R.D. Initiative, a global initiative to construct resilient data management and storage solutions. Development groups from Hitachi/Hitachi Data Systems and Oracle worked together to integrate Oracle's data integrity check algorithms into Hitachi Freedom Storage systems. Together, both companies have created technology that prevents corrupted data from being written into three principal Oracle internal file types: Oracle database files, REDO log files, and control files. Lightning 9900 V Series systems already perform extensive checksum-type validation of data integrity inside the box; therefore, DB Validator technology extends data protection to the point of entry.

## **The Collaborative Relationship**

In May 2000, the Hitachi-Oracle SANs Solution Technology Center (SSTC) was established in Japan to perform collaborative verification testing of actual customer operations, and to showcase integration and validation of advanced Hitachi Freedom Storage capabilities with the Oracle database. Technical data verified at SSTC will be translated and provided worldwide at: <http://www.sstc.gr.jp>. This will enable Hitachi/Hitachi Data Systems and Oracle to provide better and more tightly integrated solutions to the worldwide market. DB Validator technology is one of several development initiatives in which Hitachi/Hitachi Data Systems and Oracle are currently engaged.

## **Business Benefits of DB Validator and Hitachi TrueNorth™ Vision and Strategy**

### **Hitachi TrueNorth™**

DB Validator technology is part of the Hitachi Freedom Storage software portfolio. As such, it falls under the *intelligent storage systems* component of the Hitachi TrueNorth vision and strategic direction. As with all Hitachi Freedom Storage software, the DB Validator product's key value proposition is that it adds intelligence to Hitachi's superior storage technology. It accomplishes this by taking advantage of the Oracle H.A.R.D. checksum feature to ensure that corrupt data is never written to the storage system.

### **Hitachi Freedom Storage Software**

Within the Hitachi Freedom Storage software portfolio, DB Validator technology joins the Business Continuity Suite of software solutions to ensure future data availability and guarantee recoverability of data in the event an outage or loss occurs. It does this by extending the checksum-type functionality already present in Lightning 9900 V Series systems to the data's point of entry.

### **Simplify, Protect, and Optimize**

Without the timely delivery of accurate information, businesses can potentially fall prey to the competition. When data corruption occurs, the data delivery challenge is compounded because a timely and accurate resolution of the corruption is then also required.

Consider these familiar scenarios: An organization accidentally posts incorrect prices of their products on their Web site and, subsequently, loses revenue as a result of that error. Or, how many times have you tried to connect to one of your favorite Web sites and been unsuccessful because of network or technical difficulties? These types of problems exist and are compounded when corrupt data is written because

---

now two problems must be resolved: data corruption, followed by the networking component(s) error(s) that caused the corruption in the first place. The financial impact is usually greater than invalid data or network issues alone. In the case of bad data, it is usually the business owner of the data who must correct the problem, whereas, in the case of data corruption and network issues both technical and business resources are required and problem resolution usually takes more time.

With DB Validator technology, there is never any corrupt data to correct, and technical staff is notified immediately when there is an issue with the components between the Oracle databases and the Lightning 9900 V Series systems.

DB Validator technology enables resilient systems to operate 24/7 and provide the uptime required by today's most demanding enterprises. The DB Validator function protects the future availability and recoverability of data by ensuring that any corrupt data is detected immediately. This delivers a major, time-sensitive benefit: it allows users to detect and resolve problems immediately. DB Validator technology optimizes corporate resources, since issues detected immediately can be resolved more efficiently, and with a much smaller impact than issues detected after an unforeseen period of time when the data is being recovered—assuming that the corrupt data can be recovered at all. DB Validator technology also provides additional data protection by ensuring recoverability of data-type dependent elements, such as REDO logs and control file blocks.

## DB Validator Technology Details

DB Validator technology is an optional Hitachi Freedom Storage software solution for Lightning 9900 V Series systems that is designed to prevent data corruption between Oracle database platform and the storage system. This function prevents corrupted data blocks, generated in the database-to-storage system infrastructure, from being written onto the storage disk.

DB Validator technology operates in conjunction with Oracle *8i* or *9i* database applications, employing H.A.R.D. for all write data on Lightning 9900 V Series systems. DB Validator technology re-computes and validates checksums accompanying data to detect any corruption that may have inadvertently occurred as the data flowed through the many complex software and hardware layers en route from the host Oracle application to the storage system. I/O operations involving corrupt data are rejected by the system, which prevents bad data from being written to the database on Lightning 9900 V Series systems.

## DB Validator Solution Details

The DB Validator solution consists of software, a set of Fibre Channel host adapter (CHA) cards, and specific software license keys. The set of CHAs has acceleration hardware that is used to compute and validate checksums without imposing a performance penalty. The Hitachi Command Control Interface (CCI), which allows DB Validator function to be turned on and off, is provided via a license key. The license key is activated on a per-system basis, independent of the Lightning 9900 V Series system's capacity. Packaged with **Hitachi TrueCopy™** and **Hitachi ShadowImage™** software, CCI allows open systems storage administrators to create batch streams that automate operations using a command line interface. If corrupt data is detected at the port, the DB Validator function sends the I/O back to the specific instance of Oracle and permits a retry. If the second attempt to write the data fails, a "checksum" error is generated. The error is logged in **Hitachi Resource**

---

**Manager**<sup>™</sup> software, and then **Hi-Track**<sup>®</sup> initiates a “call home” to the Hitachi Data Systems support center.

## DB Validator Service Details

### Professional Services

**DB Validator Implementation Service** ensures that the installation and enablement of the solution does not impact customer operations. The service engagement includes the nondisruptive exchange of Fibre Channel adapters for existing customers, enablement of DB Validator software keys, and customer education. This service is available immediately.

### Customer Service and Support

DB Validator technology generates service information messages (SIMs) that notify Hitachi Data Systems service personnel of any corrupt data that is detected arriving at the storage system. Hitachi Data Systems responds to DB Validator generated SIMs by informing the customer that a problem exists, and then assists them with resolution on a best-efforts basis. However, since the origin of such corrupt data is external to the Lightning 9900 V Series system, resolution of the problem is beyond the scope and responsibility of Hitachi Data Systems.

All I/O operations associated with corrupt data are rejected [failed and logged in the OS log file (syslog)].

## DB Validator Functionality

### Database System and Data Corruption

As illustrated in the Figure 1, there are many components interacting between an Oracle database program and a Hitachi Freedom Storage system.

- Operating system
- I/O driver
- Logical Volume Manager (LVM)
- Fibre Channel host bus adapter (HBA)
- Fibre Channel switch (and other network components)

The many components between the database program and the storage system may fail intermittently or persistently. When these failures are detected as errors, system administrators can take appropriate recovery action. However, there still exists a very small probability that all layers might allow the data to pass without detecting the failure. The Oracle database management system and Lightning 9900 V Series systems have very robust data protection capabilities that can prevent data corruption within both database management systems (DBMSs) and the storage system. However, up to this point, data corruption that occurred between the DBMSs and storage systems could not be detected.

Another data corruption risk can occur when a non-Oracle application overwrites Oracle data files on the storage system. Lightning 9900 V Series systems and DB Validator functions also protect existing Oracle data files from being overwritten by data from other applications.

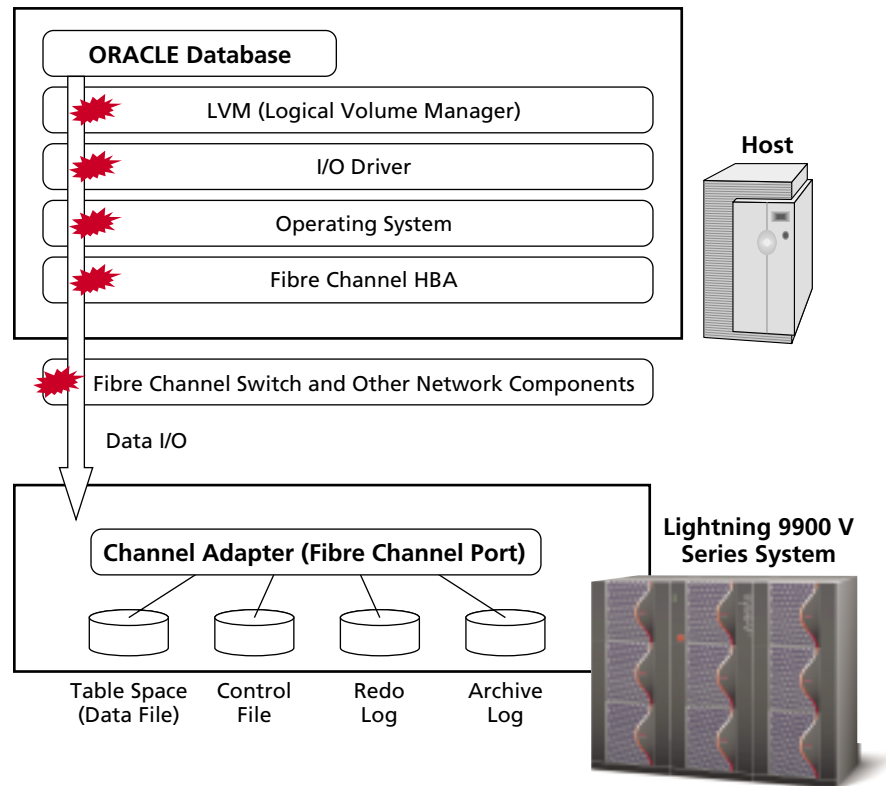


Figure 1: Sample Database Configuration.

### Validation Mechanism

An Oracle database has functionality to add the check data to each data block when issuing write requests. This check data is examined and validated at each read operation of Oracle itself. But error-detection timing is completely dependent on the read timing of Oracle databases and unpredictable to the customer.

The DB Validator function acts as a sentinel that provides additional verification of the Oracle check data on the storage system side. Figure 2 illustrates a configuration using DB Validator technology.

DB Validator technology verifies/validates the check data at *every write operation* on the storage system and reports any abnormal data *immediately* after a failure occurs. DB Validator technology will *block the incorrect write request from the hosts, return SCSI Check Condition errors, and protect the correct data in the storage system*. The Oracle database will stop itself as an error occurs, or it will separate the damaged table, depending on the error situation. Administrators can recover the database without data loss consistency because any invalid data is not written to the storage system, nor is any valid data overwritten.

DB Validator technology verifies all three file types associated with Oracle databases that are raw files or the LVM-provided raw files.

- Data files (including system table space)
- Control files
- Redo log files

The Oracle database functionality of check data is slightly enhanced with Oracle 9i Release 2. As a result, DB Validator technology has two distinctive modes, one for Oracle 8i and 9i, and another for Oracle 9i Release 2.

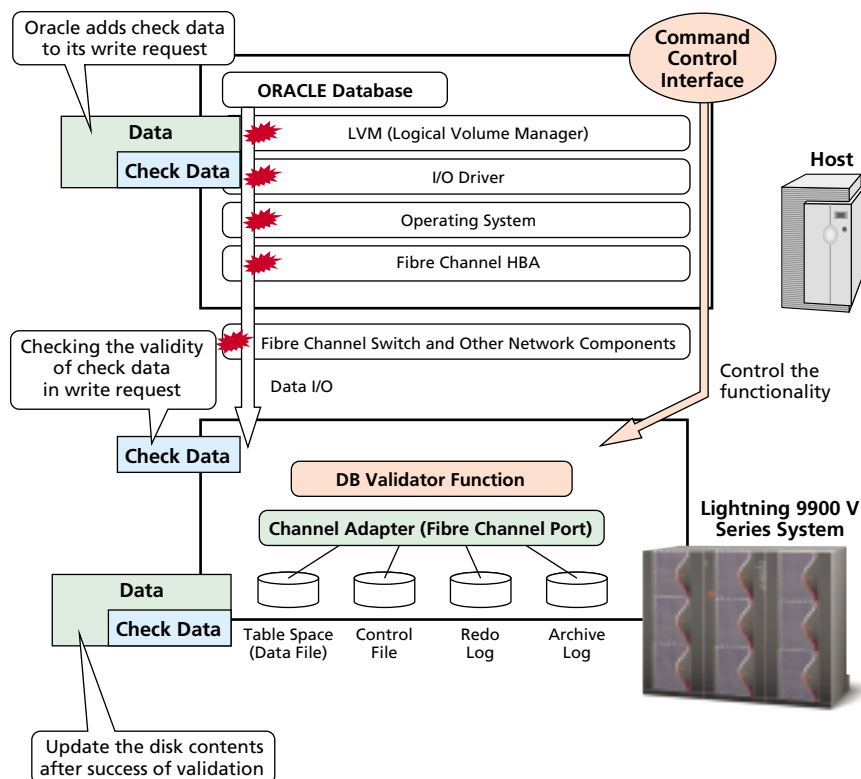


Figure 2: Sample Database Configuration with DB Validator.

CCI software is used to specify the LUNs that need to be checked and the type of verification to be performed (dependant upon the specific version of Oracle). The files to be checked include data files (including tablespace files), redo log files, and control files. The files must be stored on raw volumes. If data is written from the host server to the storage system, the Lightning 9900 V Series system's CHA initiates the data check. If invalid data is detected, the write I/O is rejected, the CHA reports the error to CCI, and CCI outputs an error to the log (syslog).

---

## Summary

The recent collaboration Oracle, Hitachi, Ltd., and Hitachi Data Systems and the resulting Hitachi implementation of the Oracle H.A.R.D. Initiative, will be of great interest to those enterprises seeking ways to ensure business continuity. Rather than address data corruption that occurs between Oracle8*i* and Oracle9*i* databases and networked storage after the fact, DB Validator technology now allows data corruption *prevention* in this area. For those business-critical applications running on an Oracle database that require the highest possible performance and availability, the best protection against data corruption is the Hitachi DB Validator function running on the most available and highest performance storage—the Hitachi Freedom Storage Lightning 9900 V Series systems. As one of the optional functions in Lightning 9900 V Series systems, DB Validator technology supports the Hitachi TrueNorth vision and strategic direction by protecting and thus helping to optimize information infrastructures.

**Hitachi Data Systems**  
**www.hds.com**  
**storage@hds.com**

**Corporate Headquarters**  
750 Central Expressway  
Santa Clara, California 95050-2627  
U.S.A.  
(408) 970-1000  
info@hds.com  
channelone@hds.com

**Asia Headquarters**  
Suite 3301-6, Shell Tower  
Times Square, 1 Matheson Street  
Causeway Bay  
Hong Kong  
2525-2385  
infoasia@hds.com

**Australia/New Zealand Headquarters**  
11-17 Khartoum Road  
North Ryde NSW 2113  
Australia  
02-9325-3300  
info@hds.com.au

**Canada Headquarters**  
2550 Victoria Square Avenue  
Suite 601  
Toronto, Ontario  
M2J 5A9  
(416) 494-4114  
www.hds.com

**Europe Headquarters**  
Sefton Park  
Stoke Poges  
Buckinghamshire SL2 4HD  
United Kingdom  
01753-618000  
info.eu@hds.com  
www.eu.hds.com

**Latin America Headquarters**  
750 Central Expressway, MS 3468  
Santa Clara, California 95050-2627  
U.S.A.  
(408) 970-1000  
infolatin@hds.com

**U.S. Headquarters**  
750 Central Expressway  
Santa Clara, California 95050-2627  
U.S.A.  
(408) 970-7447  
ussalesinfo@hds.com

Hitachi Data Systems is registered with the U.S. Patent and Trademark Office as a trademark and service mark of Hitachi, Ltd. The Hitachi Data Systems logotype is a trademark and service mark of Hitachi, Ltd. Hi-Track is a registered service mark and Freedom Storage, Lightning 9900, TrueNorth, ShadowImage, Resource Manager, and TrueCopy are trademarks of Hitachi Data Systems Corporation.

All other trade names, trademarks, and service marks used herein are the rightful property of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, express or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect, and that may be configuration-dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

©2003, Hitachi Data Systems Corporation.  
All Rights Reserved.  
WHP-127-01  
January 2003