

Hitachi Data Instance Director Software Version 5.0.1 Release Notes



About this document.....	2
Intended audience.....	2
Getting help.....	2
About this release	2
Product package contents.....	2
Upgrading the software	3
System requirements.....	3
Features and enhancements	3
Resolved problems	3
Known problems	4

About this document

This document describes the revisions that were made to Hitachi Data Instance Director, software version 5.0.0. It describes the related enhancements, any known problems and solutions and resolved issues. It supplements the main user documents for Hitachi Data Instance Director.

Intended audience

This document is intended for customers and Hitachi Data Systems partners who install, license and use Hitachi Data Instance Director.

Getting help

[Hitachi Data Systems Support Connect](https://support.hds.com/en_us/contact-us.html) is the destination for technical support of products and solutions sold by Hitachi Data Systems. To contact technical support, log on to Hitachi Data Systems Support Connect for contact information: https://support.hds.com/en_us/contact-us.html.

[Hitachi Data Systems Community](https://community.hds.com) is a global online community for HDS customers, partners, independent software vendors, employees, and prospects. It is the destination to get answers, discover insights, and make connections. Join the conversation today! Go to community.hds.com, register, and complete your profile.

About this release

This software release provides new features as described in the section Features and Enhancements. We recommend that you upgrade your software to version 5.0.1.

Software Snapshots and Hyper-V CSV support have been removed from 5.0, they will be reintroduced in a later 5 release. If using one of these features please remain on version 4.2.

As the changes in this maintenance release are limited to Oracle functionality only a Windows 64 bit version and Linux 64 bit version of HDID have been increased to version 5.0.1. It is acceptable to mix version 5.0.0 and 5.0.1 versions of HDID in one system.

Product package contents

Item	Details
Software	Hitachi Data Instance Director, version 5.0.1
Release Notes	Hitachi Data Instance Director Release Notes, RN-93HDID000-01
User Guide	Hitachi Data Instance Director User Guide, MK-93HDID001-00
Quick Start Guide	Hitachi Data Instance Director Quick Start Guide, MK-93HDID002-00

Upgrading the software

Version 4.0 onwards can be upgraded to version 5.0.1. If running an earlier version please upgrade to version 4.0 prior to upgrading to 5.0.

After you complete the upgrade installation on the master, upgrade all other nodes. To complete the upgrade, click **Save and Activate** from the tool bar to ensure that all nodes are running the correct rules.

System requirements

Refer to Chapter 2 of the *Hitachi Data Instance Director User's Guide* for information about the supported operating systems and hardware requirements.

Features and enhancements

Oracle 12c Support for Linux on Hitachi File and Block

Version 5.0.1 introduces support for Oracle 12c in the 11g compatibility mode of operation, i.e. not using the Oracle 12c introduced features.

Note that the support information in the User Guide has not been amended to reflect this change.

Resolved problems

The following previously known issues have been resolved in this release:

- 7478 If the Oracle configuration is:
 - ASM
 - Data Files on one ASM Diskgroup
 - Redo Logs on one ASM Diskgroup
 - Archive Logs on one ASM Diskgroup

The mount operation will not mount the ASM Diskgroup for the Redo Logs.

Workaround:

Manually rename the ASM Diskgroup (if an alternate name was requested by using the renaming command) and mount the ASM Diskgroup. Prior to an unmount operation unmount the ASM Diskgroup manually and rename the ASM Diskgroup back to the original name.

Known problems

- 48 When working with 64 bit Linux installations, symbolic links must be manually created in order to backup and restore ACLs and to use extended attributes with the `-s` option: `/lib/libacl.so -> /lib/libacl.so.1 /lib/libe2p.so -> /lib/libe2p.so.2`
- 70 HDID does not support SELinux in enforcing mode.
- 81 HDID uses Java on the master and repository nodes. Before updating Java on these nodes, the service should be stopped to ensure that Java updates successfully.
- 91 Currently HDID offers no protection for the "Block" operation. Blocks should only be applied to a sufficiently narrow set of classifications.
- 96 Upgrading an OS (for example, Microsoft Windows XP to Microsoft Windows Vista) may require reinstalling HDID as different components are installed for different variations of Microsoft Windows.
- 100 Files added to a zip file using the right mouse click method are not tracked correctly when monitoring only that file type. For complete access, monitor the directory.
- 101 In a forwarding topology network activity on the forwarding node is not displayed during resynchronization.
- 103 If blocking activity by users that are not in the Administrators group, Windows may report a delayed "write failed" error during certain operations. In most cases, this can be ignored because data is not lost.
- 108 When a node that has been assigned to another master is not discovered, it cannot provide any user feedback.
- 109 Moving time backwards is not recommended, as the CofioHub service will need to be restarted.
- 110 DHCP Renewal can cause a temporary disconnection of a node which will get logged as "Machine has disconnected" with a subsequent log, "Machine has connected". If a host of virtual machines creates a Windows recovery point, then the hosted virtual machines may exhibit the same issue.
- 112 It is not possible to view log attachments from the client side restore user interface.
- 117 There may be a hidden directory 'HDID-RecycleBin' remaining on the root of the drives after an uninstall. It is safe to remove this post uninstall.
- 118 A symmetric mirror can only support two nodes, however it is possible to add more than two nodes to the symmetric mover node from the user interface.
- 243 If fast resynchronization is selected, and only file metadata changes, such as file ownership or file permissions between batch backups, then the changes are not captured. These changes are only captured when the file data changes.
- 1065 After a successful DR recovery the machine may need to be restarted a second time before HDID functions correctly.

- 1309 Occasionally, a restore from a system restore point may fail. It is recommended to stop the HDID file system filter prior to the restore. From the Command Line Interface, run `sc stop dcefltr` to stop the filter.
- 2092 Installing Hitachi Data Instance Manager from the CLI or using the unattended mode will fail to install the repository capability correctly. Only use the CLI or unattended mode to install clients.
- 2158 Node names with more than 64 characters can cause problems with backup, restore, and node management functionality.
- 3079 HDID always attempts to send data directly to a node with a tape drive attached. This can cause issues with internet connected nodes that require communications to be routed through the master. If this functionality is required, then backup to a repository that is local to the tape drive, with a short retention, then forward to tape.
- 3275 It is not possible to quiesce a Hyper-V application live backup.
- 4202 If a folder is renamed during live capture of Microsoft Exchange, the event is not captured. E-mail messages will be consistent and all grouped in the same folder including any new e-mail messages added to the folder. However, a restore will bring back the old folder name. A resynchronization will resolve this issue.
- 4301 E-mail can't be exported to the file system if the resulting path is greater than 256 characters. Changing the target path for the restore may resolve the issue.
- 4420 When an Exchange User's mailbox is deleted there will be warnings logged indicating the mailbox can't be accessed. After a maximum of 12 hours the warnings will stop.
- 4534 On Windows 7 32 bit and Windows 8 the HDID installer will not make a firewall exception for HDID. This will need to be added by hand. The TCP port 30304 is required to be open for incoming traffic.
- 4628 If a namespace on an HCP exceeds its quota then the repository will stop sending data to all namespaces until the quota issue is resolved.
- 4660 When tiering to HCP, the option, fast incremental based on file modification date, should be selected. Otherwise, during a batch resynchronization, all data that is tiered to the HCP will be transferred back to the repository.
- 4677 Performing live backup with stubbing, if a file is renamed after it has been tiered to HCP but before it is stubbed it will not be stubbed until it is edited again. This is also true for all files inside a directory that is renamed. This is not an issue with batch backup.
- 4705 If sparse files are tiered to the HCP they will become non-sparse and will restore non-sparse
- 4917 On some Linux distributions the avahi-daemon can interfere with HDID network communications. This may cause backups to fail. Disabling or uninstalling the avahi-daemon resolves this issue.
- 4993 If a stubbed file is accessed, the file is de-stubbed and will be resynchronized to the repository. If the file is renamed, or moved to a location

within the policy scope before the partial resynchronization occurs, it's possible that any modifications made at that time (and from that point forward) will not be captured until you resynchronize the repository store.

- 5348 When stubbing to Azure if the stubbing criteria age is OR'ed with access it will not work as expected. Only the first of the 2 criteria will be used when determining what to stub. For example 'older than a week' or 'not accessed today' will result in files older than a week being stubbed, accessed today is ignored.
- 5355 Using Exchange individual mailbox in a parent / child domain set-up when both domains contain the same named distribution group only the parent domains distribution group will be protected.
- 5466 It is not possible to validate a Hyper-V backup when using Clustered Shared Storage. It will fail with the error: "Error creating snapshot: Failed to create snapshot : (2147754771) VSS flush writes timeout"
- 5908 To perform a snapshot of an MS-SQL database all databases including the system database need to be located on the same Hitachi block device.
- 6026 When upgrading to 5.0 there maybe errors in the logs containing "Failed to load rules". This is due to HDID 4.x rules not being compatible with HDID 5.x rules. After upgrading all clients the rules should be redistributed from the master.
- 6051 If the set-up of a replication fails HDID will not clean-up. Detailed steps to perform the clean-up by hand can be found on the HDS knowledge base
- 6080 If windows backups fail with VSS errors it is possible the HDID VSS provider needs to be reinstalled. To do this perform the following:
 1. Stop cofiohub service
 2. Uninstall HDIM VSS provider with <HDID>\bin\vssprovider\uninstall-aimstorprovider.cmd
 3. Reinstall HDIM VSS provider with <HDID>\bin\vssprovider\install-aimstorprovider.cmd
 4. Start cofiohub
- 6181 If a path is excluded that is not a sub path of an included path HDID will backup more than expected to the repository / mirror
- 6335 There is currently no check to make sure the directory selected as the mount point is empty. This means it is possible to mount a snapshot inside or even over the top of another mounted volume. This should be avoided.
- 6506 To mount a snapshot to a server the server needs to have a volume currently mounted from the Hitachi Block Storage from which the snapshot originated. This enables HDID to find the correct path between the server and the storage.
- 6526 When mounting a snapshot that contains a mounted sub directory, the subdirectory will be mounted as expected. However, the volume referenced by the subdirectory will also be mounted as a separate drive. Unmount will unmount both the expected and unexpected mounts.

- 6636 When performing a snapshot based on a hardware path, e.g. by specifying the LDEV explicitly, the only valid option when mounting is 'Expose volumes to the host but do not mount them'. However the other options are available but if used will be ignored.
- 6705 When restoring a Windows Server 2012 Machine additional physical disks, aside from the one which contains the operating system, may not be brought online or given the correct drive letter automatically. This can be rectified in the disk management screen, and the original drive letters can be found within the fixed disk snapshot information on the master which the machine was restored from.
- 6760 When selecting a target storage pool if the error 'Failed to browse: Browse timed out, storage node may be in use' is encountered retry the selection of the pool.
- 6792 When mounting a snapshot there is always only one path to the disk. To make the mounted snapshot multipath requires additional steps outside of HDID.
- 6887 Excluding a VMWare ESX virtual machine will result in all but the excluded machine being backed up regardless of the inclusions.
- 7006 If a replication needs to change from batch to live (or vice versa) the replication would need to be removed from the data flow and the rules redistributed, then the new replication defined.
- 7051 The destination of a HNAS replication pair should be treated as read-only. If changes are made to the destination a replication refresh may not make the source and destination identical.
- 7072 On upgrade from HDIM to HDID the old shortcut will be left behind, this should be deleted as it no longer functions.
- 7122 An ISM node should not be rebooted or shutdown when replications are being setup. If this happens the replications will continue, however HDID will not be aware of these replications when restarting and will attempt to set them up again. See Knowledge base for how to recover from this.
- 7163 Authorizing multiple ISM nodes from the node management screen in one operation will fail. Trying the failed authorizations one at a time will succeed.
- 7174 The current activity report is limited to repository activity only.
- 7215 It's possible when defining an Hitachi block storage agentless node to enter an LDEV range not supported by the underlying hardware, this is not checked and can result in failed backups.
- 7231 In the case of application or path classification; if the source machine is unavailable it is not possible to revert a snapshot using HDID. In the case of a hardware path classification the source machine does not need to be available but it is important to unmount the volume before performing the revert.
- 7293 Universal Replicator journals (primary and secondary) used for HDID orchestrated replication should be exclusively used by HDID.
- 7326 Revert is not a supported operation in an Exchange DAG environment
- 7369 The PDF document "Installation of Outlook form" incorrectly references the outlook form to be installed as hdim_stub_form.fdm. This has been renamed

to "HDID Exchange Email Destubber Outlook Form.fdm" The user guide correctly references the new name.

- 7409 When mounting a Linux snapshot to a different Linux machine; in order for the user and group names to be displayed correctly the users and groups must have the same ID's as the source.
- 7465 HDID does not support a Hitachi Block Storage LDEV with more than one partition / volume.
- 7522 Reverting a Thin Image snapshot where its parent is the destination of an active Remote Copy is not possible and will result in an error of "The P-VOL is sharing the S-VOL of Remote Copy" To revert the Thin Image or ShadowImage the Remote Copy relationship must first be torn down by removing it from the dataflow and redistributing the rules.
- 7538 Only LTO drives are supported when using a virtual tape library presented by Hitachi Protection Platform.
- 7591 Software snapshot support has been removed in 5.0. It will be reinstated in a future release.
- 7594 HDID installations should first be upgraded to 4.2 before upgrading to 5.0. Upgrading from versions prior to 4.2 is not supported.
- 7609 If the following error occurs during an upgrade please wait and retry, the binary is probably being slow to shut down. If the problem persists seek assistance from HDS support.

```
Error copying file from C:/Program Files/Hitachi/HDIM/xxxxxxx to C:/Program
Files/Hitachi/HDIM/xxxxxxx
```

```
-----
```

```
Abort Retry Ignore
```

```
-----
```

- 7615 If a mount operation needs to mount multiple disks and one of the mounts fails the snapshot / replication destination is shown as mounted rather than partially mounted. The logs will indicate the mount was only partially successful.
- 7633 Occasionally a push upgrade will fail. To resolve this a manual install on the node will be required.
- 7643 When defining a policy for Oracle RAC a schedule must be used. Using the 'Any-time' schedule and relying on the RPO setting only may cause more snapshots than expected. Additionally Oracle RAC nodes must have their times synchronized.
- 7667 When installing HDID on a system using Chinese locale the installer automatically uses the Chinese language, however there are missing strings. The English installer should be used by adding '--installer-language en' to the command line when running the installer. Once installed the user interface will work correctly in Chinese.
- 7671 Removing a replication from the data flow screen while the destination of the replication is mounted will cause the UI to stop showing it is mounted. A

restart of the services on the controlling ISM will resolve this issue and will allow an unmount.

- 7673 Before issuing a revert operation the user should ensure the volume being reverted is not mounted, this can be confirmed from the storage screen, managed storage tab. If it is mounted the revert will fail silently.
- 7692 Changing the configuration of a replication by removing a volume can cause an error "Handler call failed: Could not start HORCM instance" This is a transitory issue only and the system will recover.
- 7693 For mounting a hardware based snapshot to an ESX guest please see additional instruction on the knowledge base. Additional configuration of the ESX host is required for the guest to detect the disk. Until these additional steps are completed the guest will wait forever for the presented disk to appear.
- 7702 Hyper-V Clustered Shared Volume support is not available in version 5.0, for support please continue to use version 4.2. Hyper-V CSV support will be added to a future version of HDID v5.
- 7703 HNAS is not supported in a multi tenanted set-up.