Hitachi Data Instance Manager Software Version 4.2.3 Release Notes

Contents

RN-93HDIM000-10

About this document

This document describes the revisions that were made to Hitachi Data Instance Manager, software version 4.2.3. It describes the related enhancements, any known problems and solutions, and resolved issues. It supplements the main user documents for Hitachi Data Instance Manager.

Intended audience

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

Getting help

The Hitachi Data Systems Support Center staff is available 24 hours a day, seven days a week. To reach us, please visit the support Web site for current telephone numbers and other contact information:

http://www.hds.com/services/support/

If you purchased this product from an authorized HDS reseller, contact that reseller for support.

About this release

This software release provides fixes to previously reported known issues, and it introduces important new features. We recommend that you upgrade your software to version 4.2.3.

Important: There is no Linux installer for version 4.2.x. If you have Linux master or repository nodes, do not upgrade to version 4.2.x. However, if you only have Linux sources, then you can mix a version 4.2.x master or repository with version 4.1.x Linux sources.

Product package contents

Item	Details
Software	Hitachi Data Instance Manager, version 4.2.3
Release Notes	Hitachi Data Instance Manager Release Notes, RN-93HDIM000-10
Documentation	Hitachi Data Instance Manager User's Guide, MK-93HDIM001-04 Hitachi Data Instance Manager Quick Start Guide, MK-93HDIM002-02

Upgrading the software

It is important to upgrade ALL internet connected nodes first (including internet connected repositories), then upgrade the Master followed by the Repository nodes and finally 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 Manager User's Guide* for information about the supported operating systems and hardware requirements.

Features and enhancements

The following updates that are available in Data Instance Manager version 4.2.0 continue to apply to version 4.2.3. For additional information about these enhancements, refer to the *Hitachi Data Instance Manager User's Guide*:

- Data Instance Manager supports backing up file system files to Microsoft® Azure.
- The Data Instance Manager software installation includes a built-in evaluation key for sites that want to use the product on a trial-to-permanent basis. After the software is installed, it can be used for 45 days before a permanent license key is required. (Existing, purchased customers who upgrade to version 4.2.x are unaffected by the feature: Their permanent license continues to convey during the upgrade.)
- A Data Instance Manager installation continues to be supported on a 32 bit Windows OS, and with version 4.2.0, a native Data Instance Manager installation is supported on a 64 bit Windows OS.
- An executable file, **repomaint.exe**, can be used to reduce the size of the repository and free empty disk space that accumulates; for example, after a large set of data has been manually deleted, a number of snapshots have expired, or a number of stores became inactive and were deleted. It should be noted however that this process required the repository to be unmounted and can take a long time for a large repository.
- If needed, you can optionally specify the files and directories that you want exclude from being stubbed/removed, and you can block specific programs from reading and unstubbing stubbed files. A configuration file, **sourcestubdaemon.cfg**, lets you create such exclusions as separate blacklists.
- Data Instance Manager supports the use of standalone tape drives for sites that prefer not to use a tape library.
- VMWare ESX version 5.0 (and later) is supported. With Data Instance Manager version 4.2.0, support for VMWare ESX backups with 32 bit proxy nodes is discontinued.
- Microsoft Windows Server 2012 supports Hyper-V hosted on Clustered Shared

Storage. Following are the usage requirements for backing up data using Data Instance Manager:

- The Hyper-V for CSV configuration must use the **Batch** data mover operation.
- **Quiesce configured applications** must be enabled for the Hyper-V for CSV configuration.

The following features and enhancements have been added to this release:

- 6840 Diagdata can now run with or without taking snapshots. By default diagdata will snapshot the local drives; this enables logs to be captured successfully without having to stop and restart the services.

  For New installations an item "Snapshots Allowed" has been added to the diagdata.cfg and by default is set to **True**. This means diagdata will snapshot each time diagdata is run.

  When running diagdata adding the parameter --***nosnapshots*** can be used to switch snapshots off for the diagdata being run and also --***withsnapshots*** can be used to snapshot where the configuration file is set to **FALSE**.
- 6663 Logs added to capture errors for VixDiskLib. This will enable easier problem solving of VMware issues.
- HDIM now uses Microsoft Exchange Autodiscover to locate users mailbox's for HDIM mailbox backup to HCP. Because of this, Autodiscover is ***required*** to be configured in order for this to work.
- 6993 An error (instead of a warning) is now produced for quiesced application backups where the VSS Writer is missing from the source machine. This will allow a clear notification to the user that the backup was not successful and attention is needed.

Resolved problems

- 6841 Pre/Post script Log failure messages have improved detail.
- 6540 Upgrade leaves old VMware binaries in the \bin directory. All VMware binaries are now deleted from the \bin directory so that they are only contained in the \bin\VMware\vddk\bin and will no longer cause conflicts.
- 6963 Local and Central Controllers exit unexpectedly. The GUI hangs where the local and central controllers have exited. A fix has been applied to stop the Central controller exiting.
- 6932 Non-fatal error logged when backing up ESXi server. The backup completed successfully. Error no longer reported.
- 6566 VMware snapshots not always deleted following a VMware backup. Value changed to explicitly consolidate snapshots.
- 6912 Non-fatal error logged when backing up ESXi server. Error not now displayed as only relevant for VCenter.
- 7029 Resync fails due to files with certain special characters having previously being removed from the snapshot.
- 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."
- 7063 HDIM cannot handle HNAS time query failures effectively.

  When the time cannot be retrieved from a CIFS share to accurately determine which files have changed the time on the proxy node will be used in combination with a buffer of thirty seconds.

  This buffer can be altered in the resyncsender.cfg with a tag named 'CIFSTime' which is of type uint32 and represents the amount of seconds to buffer.

  If snapshots exist on the repository which have experienced this error then after an upgrade the user will need to complete another full backup so that the correct timestamp is stored within the repository.

Known problems

- 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
- Data Instance Manager does not support SELinux in enforcing mode.
- Data Instance Manager 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.

- Currently, Data Instance Manager offers no protection for the "Block" operation. Blocks should only be applied to a sufficiently narrow set of classifications.
- The Backup / Replication of encrypted files is not supported.
- Upgrading an OS (for example, Microsoft Windows XP to Microsoft Windows Vista) may require reinstalling Data Instance Manager as different components are installed for different variations of Microsoft Windows.
- 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.
- In a forwarding topology network, activity on the forwarding node is not displayed during resynchronization.
- 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.
- When a node that has been assigned to another master is not discovered, it cannot provide any user feedback.
- Moving time backwards is not recommended, as the CofioHub service will need to be restarted.
- 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.
- It is not possible to view log attachments from the client side restore user interface.
- There may be a hidden directory 'HDIM-RecycleBin' remaining on the root of the drives after an uninstall.
- 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.
- 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.
- Occasionally, a restore from a system restore point may fail. It is recommended to stop the Hitachi Data Instance Manager file system filter prior to the restore. From the Command Line Interface, run sc stop dcefltr to stop the filter.
- 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.

- Hitachi Data Instance Manager 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.
- It is not possible to quiesce a Hyper-V application live backup.
- An ESXi restore using the clone option will fail if the original VM is using a distributed switch. This is because the clone uses the same port.
- 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.
- 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.
- If a namespace on an HCP node exceeds its quota, then the repository stops sending data to all namespaces until the quota issue is resolved.
- 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.
- If sparse files are tiered to the HCP, they will become non-sparse and will restore non-sparse.
- If a file is both stubbed and de-stubbed before the next batch backup, then it will not be scheduled for stubbing again until the next time it is modified.
- 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.
- On the following Microsoft Windows operating systems, the Data Instance Manager installer will not create a firewall exception for Data Instance Manager:
  - Windows 7 (32 bit)
  - Windows 8

  Instead, the firewall exception must be specified, manually. Keep the TCP port 30304 open for incoming traffic.
- On some Linux distributions, the avahi-daemon can interfere with Data Instance Manager network communications, which might cause backups to fail. Disabling or uninstalling the avahi-daemon resolves this issue.

- 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 after you resynchronize the repository store.
- All Mailbox servers that are required to be backed up using Microsoft Exchange's individual mailbox policy, must have the Client access role installed. The servers do not need to be a part of an array, if one is present. If the Mailbox servers are members of a Database Availability Group, then all member servers must use the Client Access role in order to ensure the continuation of backups in the event of a failover.
- When using the Microsoft Exchange Individual mailbox data classification within a parent/child domain (where both domains contain the same named Distribution Group), only the parent domain's Distribution Group is protected.
- When stubbing file system files to Microsoft Azure, and using the 'OR' criterion for the stubbing age for file access, the result is not compliant with the criterion: Only the first two criteria are used for determining the stubbing requirements. For example, 'older than a week' OR 'not accessed today' will result in files older than a week being stubbed, but the 'not accessed today' requirement is ignored.