



NEON Eclipse Backup and Recovery Utilities

Eclipse iChange[®]

Eclipse iCopy[®]

Eclipse iRecover[®]

product overview

Product Overview	3
Features and Benefits	3
Powerful Choices	3
About the NEON Eclipse Backup and Recovery Utilities	4
Eclipse iChange	4
Change Accumulation Types	5
Disaster Recovery Support	5
Eclipse iCopy	6
Eclipse iRecover	7
Recovery Types	8
Incremental Image Copy Types	9
Input Sources	9
Concurrent Processes	9
About NEON Enterprise Software Products	10
About NEON Enterprise Software	11



product overview

Product Overview

NEON Enterprise Software offers a fully integrated solution that allows you to back up and recover IMS databases easily and effectively. The complicated recovery process is reduced to a set of easily learned options that require little DBA experience to master. The solution offers support for all varieties of IMS databases, and includes full disaster recovery support.

The NEON Eclipse Backup and Recovery Utilities include the following products:

Eclipse iChange A flexible, high-performance change accumulation (CA) utility that produces both standard and disaster recovery CA data sets.

Eclipse iCopy A high-speed copy utility for consistent and concurrent copies of full function and Fast Path databases.

Eclipse iRecover A reliable, high-performance recovery utility for all types of IMS databases.

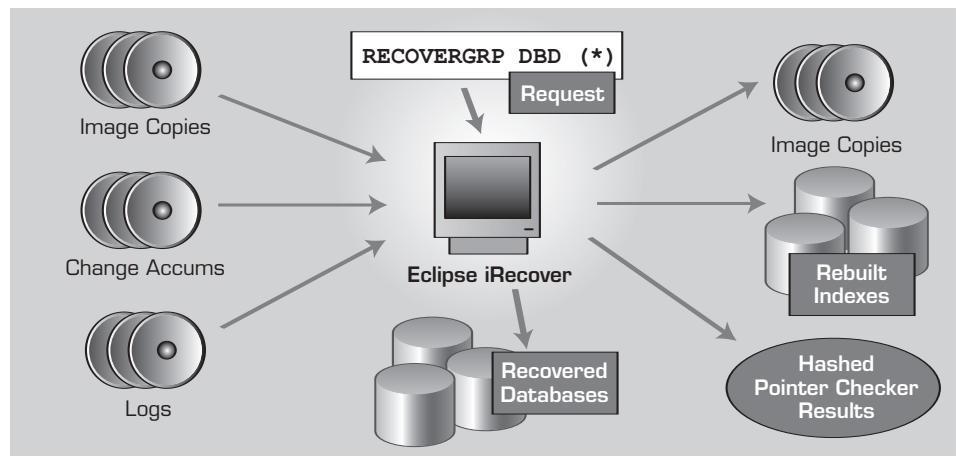


Figure 1. Recover all databases registered in the RECON data sets with a single execution.

Features and Benefits

The NEON Eclipse Backup and Recovery Utilities provide a full-featured alternative to native IMS utilities. Key highlights include:

- Provides a single toolset for recovering, accumulating changes and image copying IMS full function, Fast Path, PDF and HALDB databases
- Supports a variety of input types for change accumulation and recovery
- Allows change accumulation and recovery by DBD name, database data set, volume, DBRC grouping or database relationships
- Support for time-stamp, point-in-time, and full database recoveries, in addition to detailed support for disaster recovery operations.
- Disaster recovery support includes a load-and-go RECON backup data set and a media inventory report that details the resources needed for recovery.
- Efficient use of IMS log data sets during change accumulation and recovery processing.

product overview



About the NEON Eclipse Backup and Recovery Utilities

The NEON Eclipse Backup and Recovery Utilities are a complete set of database management tools. Each of the products in the NEON Eclipse Backup and Recovery Utilities is functionally rich and extremely flexible, allowing you to tailor each product to meet your business requirements.

Eclipse iChange

The Eclipse iChange utility is a high-performance, functional replacement for the IMS Database Change Accumulation utility (DFSUCUM0). It provides support for both standard and disaster recovery CA data sets. Eclipse iChange has the following features and functions:

- Processes multiple Change Accumulation Groups in parallel with a single pass of all required log data
- Reads all IMS log data sets in parallel
- Employs extensive data reduction and encoding to reduce the amount of data that needs to be sorted
- Performs parallel SORT processing with multiple subordinate address spaces for performance and efficiency
- Determines and allocates all necessary input data sets dynamically
- Allocates all output data sets dynamically
- Creates a report that lists all required input data sets in an easy to process format
- Integrates with the NEON Mission Control product, allowing you to view and print Eclipse iChange reports from the Mission Control repository
- Creates special disaster recovery change accumulation data sets
- Creates a load-and-go RECON backup for disaster recovery
- Creates a disaster recovery report listing all data sets required to perform recovery at a disaster recovery site, utilizing an easy to process format
- Allows disaster recovery points to be established without quiescing database activity
- Uses MVS multitasking to utilize multi-engine processor effectively
- Uses dataspace for internal queuing and processing

Change Accumulation Types

Eclipse iChange creates both normal and disaster recovery change accumulation data sets. Either or both types can be selected for a Change Accumulation Group.

Normal Normal change accumulation data sets are completely IMS (IBM) compatible and can be utilized by any recovery utility or process that supports standard IMS change accumulation data sets. Eclipse iChange produces normal change accumulation data sets that either include all available log data sets or that include only log data sets up to a specified time stamp.

Disaster Recovery Disaster recovery change accumulation data sets are not necessarily IMS (IBM) compatible. A disaster recovery change accumulation data set allows the database data sets contained within the Change Accumulation Group to be recovered to a specific time stamp, with transaction consistency, using only image copy and change accumulation data set input. The time stamp is specified when the disaster recovery change accumulation data set is created.

If no uncommitted (in-flight) updates existed at the time stamp, then the resultant disaster recovery change accumulation data set will be IMS (IBM) compatible. If uncommitted updates did exist at the time stamp, the resultant disaster recovery change accumulation data set will contain not only normal change accumulation records, but also additional records containing information necessary to restore transactional consistency when the database data sets are recovered. A recovery utility capable of understanding and processing these additional records, such as Eclipse iRecover, will be required to perform the database recoveries.

Disaster Recovery Support

The Eclipse iChange disaster recovery support is optimized for an environment where:

- No more than a few disaster recovery points are required in order to satisfy the disaster recovery data loss requirements of the enterprise. For example, the enterprise's business requirements could be satisfied by recovering to the prior 2:00 A.M.
- Database availability must not be compromised in order to establish a disaster recovery point. Using the prior example, this means that the databases are not to be /DBRed or /DBDed at 2:00 A.M.
- If data sharing is in use, the level of data sharing must not be reduced in order to establish a disaster recovery point.
- Data base disaster recovery operations are to be performed using image copy and change accumulation data sets. Log data is not to be used and, therefore, is not sent off-site.

In order to support the above environment, Eclipse iChange provides a number of functions above and beyond the creation of special disaster recovery change accumulation data sets. These additional functions are intended to assist in the management of the disaster recovery process.

Disaster Recovery RECON Backup The disaster recovery RECON backup is a load-and-go RECON backup. The RECON backup is restored at the disaster recovery site and used for disaster recovery operations and subsequent production activities. Ready to use, no additional cleanup operations are required for the backup. It is in a format that supports large RECON record sizes (records greater than 32K in length).

RECON Restore Program Eclipse iChange provides an easy to use, restore program for use at the disaster recovery site.

Disaster Recovery Media Inventory Report This report lists all data sets required to perform disaster recovery operations for the database data sets that are members of the processed Change Accumulation Groups. It includes not only the image copy and disaster recovery change accumulation data sets, but also the associated RECON backup data set.

With an easy to process format, the report provides the information necessary to manage the disaster recovery process and ensure that the necessary data sets have been sent to the disaster recovery site before a disaster occurs.

Eclipse iCopy

Eclipse iCopy is a functional replacement of the IMS Image Copy utility (DFSUDMPO), bringing the database image-copy process up to the standard expected of a high performance, state-of-the-art utility program. Eclipse iCopy provides the following features and functionality:

- Creates either consistent or concurrent image copies.
- Stacks image copies on a single tape drive.
- Creates multiple copies of an image copy.
- JCL and control-card compatible with the IMS Image Copy, with the exception of STEPLIB.
- Self-tuning and optimizing so you do not need to be concerned with items like block sizes, buffer pool sizes, or I/O subsystems.
- Dynamically allocates the data sets using DFSMDA members from STEPLIB or IMSDALIB.
- Much faster than IMS Image Copy, where n is the number of volumes over which a database data set is equally spread. Eclipse iCopy is even faster than the IMS Image Copy when the database data set is only on a single volume.
- Has the capability to provide unload, pointer checking, scanning, index building and reloading of a database in the same job that the image copy is made. The desired output is produced from a single pass of the database.
- Creates image copies that are fully compatible with Database Recovery Control (DBRC) and all other IMS recovery tools.
- Accepts utility control statements used by the IBM Image Copy utility as well as the free-form NEON Eclipse Utilities command language.
- Synchronizes volumes of multiple image copy data sets made to tape by ensuring each relative volume in the data set contains exactly the same number of records. Synchronized volumes can be useful when an I/O error occurs reading one of the data set copies.

Eclipse iRecover

The Eclipse iRecover utility is a high-performance, functional replacement for the IMS database recovery utility (DFSURDB0) and various other IBM/OEM recovery and image copy products. Eclipse iRecover provides the following recovery capabilities:

- Allows high-level recovery requests by:
 - Database (specific or generic DBD names)
 - Database data set or AREA
 - Change accumulation, DBDS, or DB group
 - Volume serial number
- Recovers related database data sets and AREAs
- Recovers database data sets and AREAs only if recovery is actually needed
- Supports full, time stamp, and point-in-time recovery types
- Provides true point-in-time backout support for full function and HALDB databases
- Creates shadow copies of IMS Full Function database data sets
- Creates shadow copies of IMS Fast Path AREA data sets
- Restores shadow copies of IMS Full Function database data sets and Fast Path AREA data sets from “archival” image copies
- Accepts recovery input from various sources
- Supports and executes several concurrent processes
- Determines and allocates all necessary recovery input data sets dynamically
- Reads all IMS log data sets in parallel
- Reads change accumulation data sets in parallel with any IMS log data sets
- Produces image copy data sets for recovered objects (optional)
- Rebuilds indexes meeting source = target criteria (optional)
- Uses MVS multitasking to utilize multi-engine processors effectively
- Performs hashed pointer checking for recovered objects
- Allocates all output data sets dynamically
- Recovers multiple database data sets/AREAS in parallel, whenever possible

Eclipse iRecover also provides these incremental image copy capabilities:

- Allows high-level incremental image copy requests by:
 - Database (specific or generic DBD names)
 - Database data set or AREA
 - Change accumulation, DBDS, or DB group
 - Volume serial number
- Makes incremental image copies for related database data sets and AREAs
- Accepts input from various sources
- Supports and executes multiple concurrent processes
- Determines and allocates all necessary input data sets dynamically
- Reads IMS log data sets in parallel

- Reads change accumulation data sets in parallel with any IMS log data sets
- Performs hashed pointer checking during the image copy process
- Allocates all output data sets dynamically
- Creates image copies of multiple database data sets/AREAs in parallel whenever possible

(Because image copy stacks (tape volumes) cannot be shared, Eclipse iRecover recovers, and makes incremental image copies of, database data sets/AREAs with image copies stacked on the same tape volumes serially.)

Eclipse iRecover provides the following system functions:

- Creates a fully conditioned RECON backup data set for disaster recovery, with supporting cleanup operations and reports
- Lists application restart symbolic checkpoint ids and in-doubt DB2 UOR status at the user-specified time stamp

Recovery Types

The Eclipse iRecover utility performs or assists the following types of recovery:

Full Recovery	Recovers database or AREA data sets to their most current state.
Time-Stamp Recovery	Recovers database or AREA data sets to a prior state, following the current DBRC rules for a time stamp recovery. Because of restrictions imposed by DBRC (no open ALLOC records), Eclipse iRecover does not include uncommitted updates in this type of recovery.
Point-In-Time Recovery	Recovers database or AREA data sets to a prior state, without regard for DBRC rules (recovery time can be within the scope of an open ALLOC record) or for the existence of a quiet point (time when no IMS subsystems have “in-flight” updates against the databases). Eclipse iRecover maintains transactional consistency when it performs a point-in-time recovery. The recovery includes only those updates that were committed up to the specified recovery time stamp.
Disaster Recovery	Recovers data after catastrophic system failure, often from resources that are stored at a separate location. Disaster recovery may involve one or more database recovery types. Eclipse iRecover helps you to determine the appropriate time stamp for a disaster recovery and creates a RECON backup that you can maintain as a disaster recovery resource.

Incremental Image Copy Types

Eclipse iRecover creates two types of incremental image copies:

Batch Image Copy	Produced when no image copy time stamp is provided, the image copy time stamp corresponds to a valid DBRC time stamp recovery point (no open ALLOC records), or the incremental image copy request is for a shadow copy of the Full Function database data set. The image copy maintains transactional consistency.
Concurrent Image Copy	Produced when the image copy time stamp provided is within the scope of an open DBRC ALLOC record and the image copy request is for the registered copy of the database or AREA data set. Unlike standard IMS concurrent image copies, transactional consistency is maintained.

Eclipse iRecover determines which type of image copy to produce based on the incremental image copy time stamp (optional) and whether the requested image copy is for the registered copy of the database or AREA data set, or for a shadow copy of the database data set.

Input Sources

Eclipse iRecover accepts recovery/incremental image copy input from the following sources:

- Standard IMS-compatible batch, concurrent, or online image copy data sets (stacked or not stacked)
- Compressed (BMC Level 1 FSE compatible) batch, concurrent, or online image copy data sets (stacked or not stacked)
- Standard IMS-compatible change accumulation data sets
- Eclipse iChange change accumulation data sets
- IMS log data sets

Eclipse iRecover accepts control statement input from the following sources:

- Eclipse iRecover native control statements
- IMS DBRC GENJCL.RECOV commands

Concurrent Processes

Eclipse iRecover is designed and optimized to recover or make incremental image copies of multiple database data sets or AREAs. The utility also processes high-level requests that could result in the recovery or incremental image copy of hundreds or thousands of database data sets or AREAs in a single operation. In fact, you can use a single control statement to request a recovery or an image copy of all registered database data sets and AREAs.

Because you do not have an infinite amount of hardware resource (tape drives, real memory, and virtual memory), Eclipse iRecover allows you to specify how many concurrent processes you want to support. For example, you may want to control the following concurrent processes:

- Log input tasks (input tape drives)
- Index rebuild tasks (concurrent SORTS)
- Database data set or AREA recovery tasks (input/output tape drives for Image Copies)

Eclipse iRecover executes all of these processes concurrently, and each process class or type may consist of multiple servers or tasks. You can specify the number of tasks or servers for each process class or type as an installation option.

About NEON Enterprise Software Products

NEON Enterprise Software offers a variety of solutions to increase and maintain data availability for your mainframe enterprise. Every NEON Enterprise Software solution is architected to work smarter than other offerings, not just faster, providing the highest levels of control and availability for your applications and infrastructure.

Database Director™

Database Director enables online reorganizations of all types of IMS full function databases without requiring an application outage. In addition to the Database Director Batch product described in this document, Database Director Online and Database Director Persist products provide maximum data availability for online databases. While reorganizing an online database, Database Director can also perform other database tasks such as cloning, space allocation, tuning, restructuring, and facilitating hardware data compression.

Eclipse iExtract™

Eclipse iExtract is a powerful utility that quickly and efficiently extracts data from both IMS full-function and Fast Path databases. Because Eclipse iExtract directly accesses the database, its performance is unmatched.

Eclipse iLM™

Eclipse iLM provides an affordable, comprehensive set of tools for cleaning and maintaining IMS and CICS libraries, including ACB, DBD, PSB, and dynamic allocation libraries, DBRC, and the DFSDDIR member of MODBLKS. By verifying that IMS-related libraries are in sync with one another, Eclipse iLM ensures database integrity and availability.

Eclipse iRepair™

Eclipse iRepair is a powerful tool for viewing, analyzing and repairing IMS database data sets and other z/OS data sets. You can use iRepair to resolve pointer check errors or other types of data errors, reducing the amount of maintenance required to back out and restore problem database data sets.

Eclipse Reorganization Utilities™

The Eclipse Reorganization Utilities are the fastest IMS reorganization utilities available. These IMS database utilities include Eclipse iBuild, Eclipse iCheck, Eclipse iCopy, Eclipse iLoad, Eclipse iSurvey, and Eclipse iUnload, all of which can be used standalone or as an integrated solution.

HALO™

HALO is a powerful new solution that provides near-continuous availability for IMS database partitioning and other restructuring. Online outages are reduced from hours to just seconds, allowing you to partition or restructure even the most critical databases without suffering long application outages that affect your business.

Lightning Utilities

The Lightning Utilities offer a streamlined, effective solution for IMS Fast Path DEDB database capacity and performance. They include Lightning Extend Online, Lightning Extend Instant, and Lightning DEDB.

Mission Control™

Mission Control is an intelligent IMS data management console that allows you to monitor and control all of the IMS full function and Fast Path databases in your enterprise. Mission Control automates database monitoring and problem resolution, enabling service-level agreements to be easily met.

Online Reorganization Director™

If you are using IBM HP utilities, Online Reorganization Director provides 100% application availability during reorganizations—plus seamless integration with IMS High Performance utilities.

Partitioned Database Facility™

Partitioned Database Facility increases IMS database capacity and improves database performance and availability, providing a cost-effective method for growing your business without affecting business applications.

Prefix Update™

Prefix Update performs prefix resolution and prefix update operations in a single job step, making the process faster and more efficient than with other solutions.

DB2 Products

Partnering with Software Engineering GmbH, NEON Enterprise Software presents a comprehensive set of solutions to improve and maintain DB2 database and application performance. The following products are available to serve the DB2 enterprise. To fully explore how NEON Enterprise Software DB2 products can help you better control your DB2 environment and improve database availability, visit www.neonesoft.com/db2.shtm.

iServe Managed Services

iServe managed services for IMS gives you the opportunity to extend your IMS expertise by providing needed services to your organization. To fully explore how NEON Enterprise Software can supplement your IMS staff and expertise, visit www.neonesoft.com/ISV.shtm.

About NEON Enterprise Software

NEON Enterprise Software is the technology leader in enterprise data availability software and services. In a world where every second counts, our tools maximize database performance and availability and minimize business risk. Founded in 1995, NEON Enterprise Software is headquartered in Sugar Land, Texas and serves customers worldwide with its dedicated team of industry experts.

For more information, visit www.neonesoft.com or call 281.491.6366 or 888.338.6366.

Copyright ©2006 NEON Enterprise Software, Inc. All rights reserved. Eclipse iChange, Eclipse iCheck, and Eclipse iRecover are registered trademarks of NEON Enterprise Software. Database Director, EADO, Eclipse iBuild, Eclipse iCopy, Eclipse iExtend, Eclipse iExtract, Eclipse iLM, Eclipse iLoad, Eclipse iRepair, Eclipse iSurvey, Eclipse iUnload, Eclipse Reorganization Utilities, HALO, iServe, iServe DBA, iServe SP, Lightning DEDB, Lightning Extend Instant, Lightning Extend Online, Lightning Utilities, Lightning X, Mission Control, NESS, Record Reorganizer and TITAN Archive are trademarks of NEON Enterprise Software. PDF is a trademark of NEON Systems, Inc., in the USA and in other select countries, and is licensed to NEON Enterprise Software. All other trademarks are the property of their respective owners. 6/08

