NEON Enterprise Software Capacity Management:

Partitioned Database Facility™
NEON Eclipse Reorganization Utilities™, Partitioned Support Edition

product overview

NEON
Enterprise Software, Inc.

Table of Contents

product overview

A Unique Capacity Management Solution

The NEON Enterprise Software Partitioned Database Facility (PDF™) increases IMS database capacity and improves database performance, providing a cost-effective method for growing your business without impacting business applications. PDF allows you to manage capacity growth without redesigning IMS or requiring an all-or-nothing migration to partitioned databases.

PDF ships with the NEON Eclipse Reorganization Utilities, Partitioned Support Edition (PSE) for efficient, cost-effective database maintenance. PDF performance, combined with the speed and efficiency of the utilities, provide an extremely effective, moderately priced IMS capacity management solution. When coupled with the NESI Database Director (D^2) product, PDF and the NEON Eclipse Reorganization Utilities/PSE form a complete capacity management solution that offers superior performance and uses fewer IMS resources than any other solution on the market.

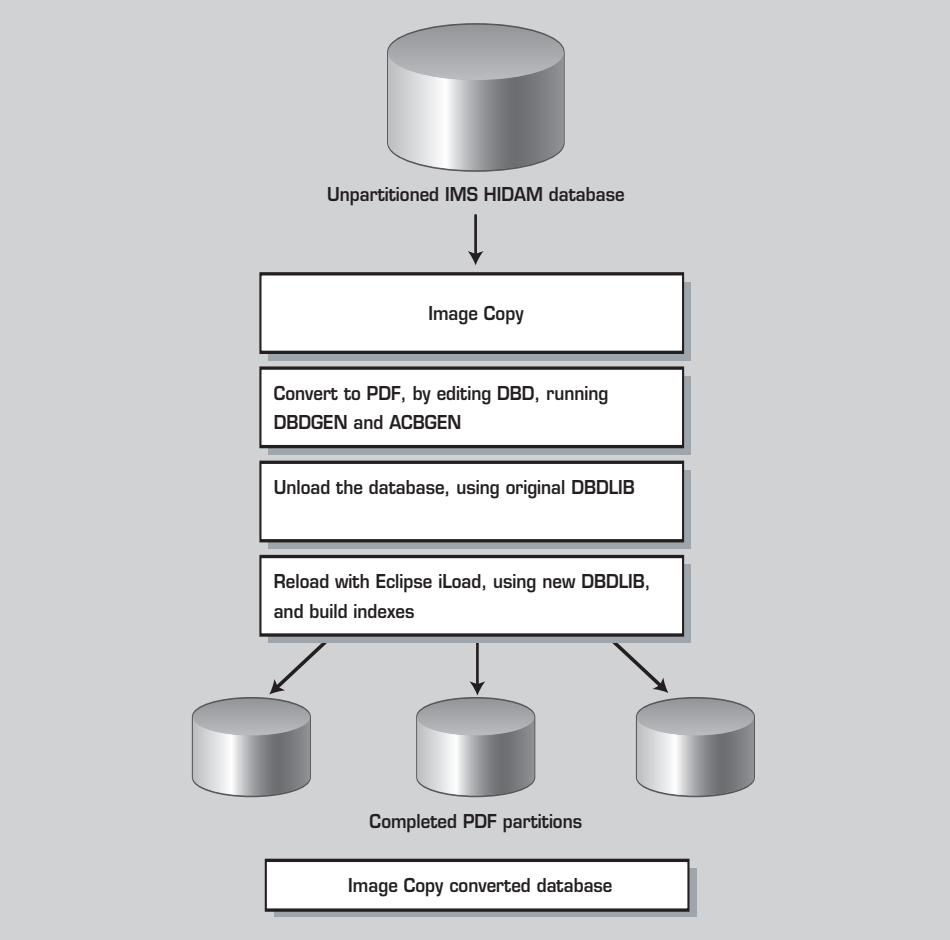

Figure 1. Partitioning with PDF

PDF Highlights

PDF allows you to organize databases into partitions and data set groups and expands database capacity, with the following features:

- Increases VSAM limits from 4 GB to more than 500 GB
- Increases OSAM limits from 8 GB to more than a terabyte
- Allows up to 127 database partitions or data set groups, expanding IMS database and index capacity by a factor of 100.

- Allows partitioning without PSB changes, minimizing the impact on IMS applications.
- Supports databases that are not registered with DBRC, reducing the learning curve for conversions in test environments.
- Supports logical relationships with unpartitioned databases.
- Allows the retention and partition of secondary indexes.
- Improves performance by reducing or eliminating data contention and I/O bottlenecks and simplifies many IMS database management tasks, such as converting databases from full function to partitioned.

NEON Eclipse Reorganization Utilities/PSE Highlights

The NEON Eclipse Reorganization Utilities/PSE provide an efficient, flexible database management solution for PDF databases. The NEON Eclipse Reorganization Utilities provide you with the power and functionality you need to reorganize and maintain even the largest and most disorganized database. The NEON Eclipse Reorganization Utilities/PSE solution set includes:

- Eclipse iUnload, a high-performance, self-tuning, unload utility that is a functional replacement for IMS HD Unload and similar unload products.
- Eclipse iCheck, a pointer-checking utility that checks database integrity and diagnoses "broken" databases.
- Eclipse iLoad, a fast utility for loading IMS databases; it can also be used with NEON iBuild to create index databases for the database being loaded.
- Eclipse iBuild, an innovative and integrated utility for creating and recreating IMS INDEX databases.
- Eclipse iSurvey, a utility for scanning and reporting information about IMS databases, areas, or partitions, including a high-keys report for databases to be partitioned.

With the NEON Eclipse Reorganization Utilities/PSE, NESI provides an optimal database maintenance suite for PDF databases. The following sections describe the benefits you can derive if you use these utilities for PDF database reorganizations and maintenance. For more information about each utility, see the NEON Eclipse Reorganization Utilities Product Overview.

Consistent Performance—Regardless of Database Organization

With other utilities, the more disorganized the database, the longer the reorganization takes. The NEON Eclipse Reorganization Utilities/PSE have consistent performance-no matter how disorganized the database. Large, disorganized databases can be reorganized at the same speed and cost as large, organized databases. You save computing time and CPU cycles, both of which reduce operating expenses and ensure your business applications are available the maximum amount of time. You also have the benefit of predictability. You can schedule your work and jobs more effectively since you can predict how long the reorganization will take.

Parallel Processing

The NEON Eclipse Reorganization Utilities/PSE can process database volumes, partitions, or areas in parallel, reducing the amount of processing time. The utilities can multiple partitions in parallel as well as perform multiple concurrent tasks. Your reorganization is completed at maximum speed, reducing your database maintenance costs and making the databases available more quickly.

Perform Multiple Operations in a Single Job Step

Since the NEON Eclipse Reorganization Utilities are an integrated solution, you can perform several operations on your databases in a single job step. For example, you can reload the database, build indexes and image copy the database in a single job step. This makes reorganizing your databases simple and efficient. Additionally, the database is only read one time, regardless of how many operations are being performed. This minimizes I/O and reduces elapsed time and CPU utilization.

High-Keys Report for Easy Partition Design

Determining how a very large database should be partitioned can be a time consuming process. Using Eclipse iSurvey, you can easily generate a high-keys report to determine the best partition strategy for the PDF database.

Capacity Management and Performance Benefits: A Look at Partitioning Solutions

The space limits inherent to IMS databases make it difficult to grow your business efficiently. PDF offers an alternative to both standard IMS databases and HALDB migrations, making it a cost-effective choice for capacity planning. The following sections describe how PDF takes advantage of IMS features to provide cost-effective capacity planning. For a concise summary of the PDF conversion process, see page 8.

Improved Data Set Support

By supporting up to 8 GB for an OSAM data set, PDF requires fewer data sets than HALDB. In addition, PDF does not create extra data sets (such as the ILDS). Creating partitions with PDF is similar to adding a data set group, rather than creating a completely new set of databases. Table 1 compares data set attributes and limits for PDF and HALDB.

Support for	PDF	HALDB
Maximum data set size	4 GB for VSAM and KSDS data sets 8 GB for OSAM data sets	4 GB for both OSAM and VSAM data sets
Number of data sets per data set group	10	10
Maximum number of partitions and data set groups	127	1001
Partition independence	NO	YES (with application changes)

Table 1. Data Sets and Partitions

Database Types and Features

PDF supports HISAM and SHISAM databases, in addition to the HDAM and HIDAM database types supported by HALDB. See Table 2 for a comparison of databases and features supported by the two solutions.

PDF supports logical relationships between partitioned and non-partitioned databases. With HALDB, you must migrate all logically related databases, or remove the relationships. Depending on the number of related databases, this can be a costly process.

Because secondary indexes are a special kind of logical relationship, they are not directly supported by HALDB or Fast Path databases. HALDB requires that secondary indexes be converted to PSINDEXes and migrated with a separate operation. The PSINDEX conversion can up to three times more space for each index.

Both HALDB and PDF require physically paired logical relationships. Virtually paired segments must be converted to physical pairs for both products.

PDF allows you to retain and partition primary and secondary indexes when you partition the database, saving you time and effort. The number of partitions in the primary index must be the same as those for the database data sets. For secondary indexes, the number of partitions can be different than the number of database partitions, providing a great deal of flexibility in handling index conversions to PDF.

There are no changes required for the secondary index DBD or IDCAMS definitions for the KSDS, if there are no shared or nonunique secondary indexes. For nonunique secondary indexes, the change can be as simple as adding an /SX field to the index to make it unique. In addition, the size of the secondary index does not increase when you convert the database to PDF partitions.

PDF supports both recoverable and nonrecoverable primary and secondary indexes. You can restore indexes from image copies and by using point-in-time recoveries.

Support for	PDF	HALDB
Secondary Indexes	YES	YES (PSINDEX)
Logical Relationships	YES	YES (only to other HALDBs)
HDAM databases	YES	YES (PHDAM)
HIDAM databases	YES	YES (PHIDAM)
Root-only HISAM databases	YES	NO
SHISAM databases	YES	NO
Virtual Paired Relationships	NO	NO
Index build	YES	NO
Index per partition	YES	YES
One index for all partitions	YES	NO

Table 2. Database Types and Features

Archiving Support

PDF allows you to continue to archive and purge inactive segments during reorganizations. You can also archive and discard aged data by updating the partition number during a reorganization. This flexibility lets you keep your current processes as your business continues to grow.

Database Maintenance and Test Environments

PDF recommends, but does not require DBRC registration for partitioned databases. HALDB requires that all partitions be registered with DBRC. In some test environments, where databases are constantly recovered or recreated, requiring DBRC can cumbersome and might require a learning curve for some DBAs and application programmers.

In addition HALDB requires PSINDEXes and ILDS data sets that are nonrecoverable and cannot be image copied by IMS, requiring procedural changes for IMS recovery. HALDB self-healing pointers can also affect database maintenance and performance.

Support for	PDF	HALDB
DBRC Required	NO	YES
Image copy indexes	YES	NO
Image copy secondary indexes	YES	NO
Indirect List Data Sets (ILDS)	NO	YES (for each HALDB partition)
Extended pointer	NO	YES
Self-healing pointers	NO	YES
Dynamic allocation	YES	NO
New IMS status codes	NO	YES

Table 3. Database Maintenance

Improved Database Performance

While any partitioning solution can improve database performance, PDF avoids the performance degradation incurred by self-healing pointers. The way that PDF handles indexes and parallel processing of partitions enhances the performance of IMS databases as it increases capacity.

Online Reorganization Performance

PDF is supported by the NESI Database Director (D^2) product for online reorganizations. D^2 provides an enormous performance improvement (both in run time and CPU used) to the online reorganization utility supplied with HALDB. D^2 reorganizations handle all IMS commands required during the process, and incur significantly less overhead than the HALDB OLR, which includes both copious log output and the time lost on self-healing pointers. You can also terminate a D^2 reorganization at any time, rolling back the operation with no impact to your system.

Support for	PDF	HALDB
Online reorganization tool	Database Director	HALDB Online Reorganizer
Clone partitions	Database Director	NO

Table 4. Online Reorganization Tools

Conversion Process Summary

Much of the cost of partitioning occurs during initial implementation. With a simplified installation and conversion process PDF reduces the overall time and cost required to implement a partitioned database solution. It typically takes 60 days or less to convert databases to PDF and have PDF databases in production.

Table 5 compares the steps required to convert an IMS database with secondary indexes to PDF or HALDB. In both cases the database must be stopped for at least some of the steps.

PDF	HALDB
1. Select database candidate to partition: HDAM, HIDAM, HISAM, SHISAM databases including primary and secondary indexes.	1. Select database candidate to partition: HDAM or HIDAM.
2. Modify DBD with PART statements.	2. Unload the database.
3. Modify DFSMDA member if necessary.	3. Modify DBD with HALDB statements and run DBDGEN.
4. For HDAM, provide number of RAAs, or for HIDAM, provide high-key values. You can use Eclipse iSurvey to generate a high-key report for the HIDAM database.	4. Delete pre-HALDB database information from RECON data sets.
5. Run DBDGEN.	5. Define partitions.
6. Run ACBGEN.	6. Allocate database data sets.
7. Image copy the original database.	7. Initialize partitions.
8. Unload the database using the original DBDLIB and the HD Unload utility.	8. Reload partitioned database as HALDB .
9. Define the DBD, DD and data set names for both the database data sets and indexes to be partitioned.	9. Image copy the database data sets.
10. In a single job, reload the PDF database with Eclipse iLoad and rebuild the indexes with Eclipse iBuild.	10. Define DBD for each secondary index and run DBDGEN
	11. Delete pre-HALDB secondary index information from the RECON data sets.
	12. Define partitions.
	13. Allocate database data sets.
	14. Sort the output file from the unload of the indexed database
	15. Load secondary index.
	16. Image copy the database data sets created from the secondary index.

Table 5. Conversion Process Summary

Making the Right Choice for Capacity Management

As data grows and database performance becomes strained, your choice of an effective capacity management solution may be the most important business decision that you make. If your IMS environment is already complex, PDF and the NEON Eclipse Reorganization Utilities/PSE can simplify the issue of capacity management by:

- Supporting your existing data structures, such as secondary indexes and logical relationships with minimal conversion requirements.
- Providing easy-to-use tools for conversion to PDF, such as partition selection routines and the Eclipse iSurvey high-keys report.
- Implementing partitioned databases in a familiar, easy-to-define method that does not require new types of data sets or additional IMS resources.
- Performing database maintenance tasks at the maximum possible speed, using parallel processing and a single job step to accomplish multiple tasks.

The following sections provide a technical overview of PDF, including the way each type of IMS database is partitioned, and the simple control statements used to define a partitioned database.

PDF Technical Overview

Converting an IMS database to PDF is similar to creating a data set group, requiring a simple modification of the IMS DBD, an ACB change, and allocation of the PDF database data sets. No changes are required to the associated PSBs, so the applications that access the database need not change.

- The DBD must be modified and a DBDGEN is required to implement the changes. At least one PART statement must be added to the DBD. You can also include control statements to specify a partition selection routine or data set group.
- An ACBGEN is required to generate a DMB for the partitioned DBD. When an ACBGEN is performed, an online change can be used to introduce the new DMB into the IMS online system without having to restart the system.
- DBRC registration for partitioned databases is the same as with data set groups. The INIT.DBDS is used to register each partition in DBRC. DBRC registration of partitioned databases is not required; however, using DBRC to keep track of image copies and generate JCL for change accumulation and recovery of partitions is extremely useful.
- Partitioned databases can by allocated using DD statements in JCL or DFSMDA dynamic allocation members. There is a single DFSMDA member for a partitioned database with the ddnames and data sets names for all partitions defined using TYPE=DATASET statements. Dynamic allocation members are recommended for partitioned databases to avoid changing JCL DD statements when adding or deleting partitions. Dynamic allocation members allow partitions to be added with on-line change.

The following sections describe the product's technical requirements and the partition methods for each type of IMS database that PDF supports.

System Requirements

The capacity management products require the following software be installed on your mainframe system:
- An IBM-supported version of IMS
- An IBM-supported version of z/OS

Product Limits

The following database features or types are not supported by the capacity management solution:

- DOS compatible indexes
- Hierarchical pointers
- Nonunique secondary indexes that use KSDS/ESDS pairs
- Shared secondary indexes
- The checkpoint/restart function of UCF database load
- XRF

Supported Environments

The capacity management solution supports partitioning for the following IMS full function database types and features:

- BATCH DBB and DL/I
- Data set groups
- Data sharing at the database and block levels
- Databases that are not registered with DBRC
- Databases with primary and secondary indexes
- DB/DC
- DBCTL
- FDBR
- HDAM, HIDAM, root-only HISAM, and SHISAM databases
- IMS utilities
- IRLM
- Logically related databases

HDAM Databases

PDF uses one of the following methods to partition HDAM databases:

A standard HDAM randomizing algorithm (a single-stage randomizer)—When using a single-stage randomizer, PDF determines the number of blocks in the entire database and passes this information to the randomizer. The randomizer examines the root segment key and selects the root segment pointer address, known as the Root Anchor Point (RAP) to determine the storage location of each root segment in the database. PDF determines where to partition the database, based on the number of blocks in the root addressable area (RAA) of each partition.

A DEDB randomizing algorithm (a two-stage randomizer)—A two-stage randomizer randomizes within a particular partition to keep all records for specific criteria (such as geographic area) within the same partition. The randomizer determines into which partition/area to place a root segment. The routine then determines the location of the root segment within that partition/area. (See Figure 2.)

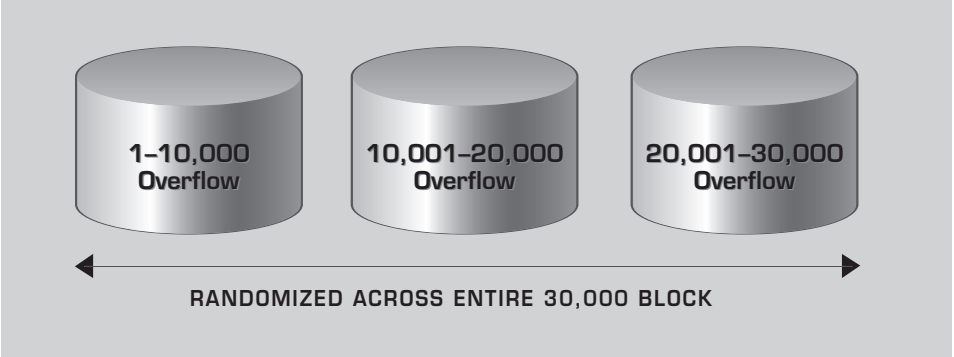

Figure 2. Partitioning an HDAM Database

Partition Selection Routine—The routine selects the partition. The partition selection routine is specified in the DBD by using a PARTDBD control statement. For more information about the PARTDBD control statement, see page 13.

Without a partition selection routine, PDF uses the name of the randomizer load module to determine which interface to use. If the load module name begins with DBF, PDF uses the DEDB randomizing algorithm. For any other load module name, PDF uses the standard HDAM randomizer.

HIDAM Databases

PDF partitions HIDAM databases by the primary index key using the following methods:
- The high key for each partition is defined, establishing a range of root keys for each partition
- A partition selection routine is used to select the partition.

See Figure 1 on page 3 for an example of a partitioning a HIDAM database. If the primary index is not partitioned, then only the data portion of the HIDAM database is partitioned.

The primary index contains the keys for all partitions. The high key is used to separate the HIDAM database into partitions. The high key in each partition is also used to position for root insert and for twin backwards pointers on the root. Each high key in a HIDAM partition ends in hex 'FF' and is inserted by IMS when the database is loaded. As with the hex 'FF' key in nonpartitioned HIDAM, the high key in each partition cannot be inserted or retrieved by the application program.

HISAM Databases and SHISAM Databases

You can partition HISAM and SHISAM (root only) databases using high keys, similar to how you partition indexes and HIDAM databases. An advantage of SHISAM databases is that they can be processed by VSAM applications outside of IMS. Any VSAM applications must be aware the database is partitioned into multiple KSDS data sets. Partitioned root-only HISAM and SHISAM databases use the support provided by partitioned indexes.

Data Set Groups

DL/1 supports multiple data sets with data set groups, but requires all segments of the same type to be in the same data set. Data set groups allow a dependent segment to be stored in a separate data set from the root. (See Figure 3.)

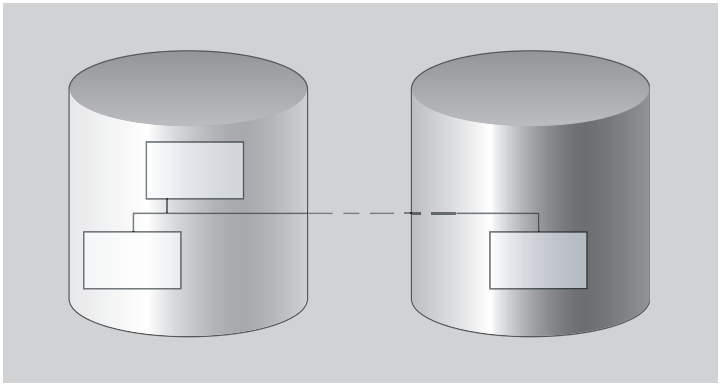

Figure 3. Creating Data Set Groups

This can be useful for large, infrequently accessed segment types. An extra I/O is needed to read the entire database record since the dependent segment is in a separate data set.

Data set groups are useful for databases greater than 4 GB when one or two segments comprise a significant percentage of the database, and the segments are not frequently accessed or updated. However, data set groups are not always a permanent solution to the problems of large databases.

Partitioned databases implement multiple database data sets similar to DEDB Areas. All segments in the same physical database record are in the same partitioned data set. With both HDAM and HIDAM, the partitioning is done by the root key. All physical dependents are stored in the same data set as the root. (See Figure 4.) A database that is partitioned uses at least one, primary data set group. A partitioned database can have more than one data set group, but no more than 10 data set groups. The number of data set groups multiplied by the number of partitions must not exceed 127. For example, if you designate 31 partitions, you can have no more than four data set groups.

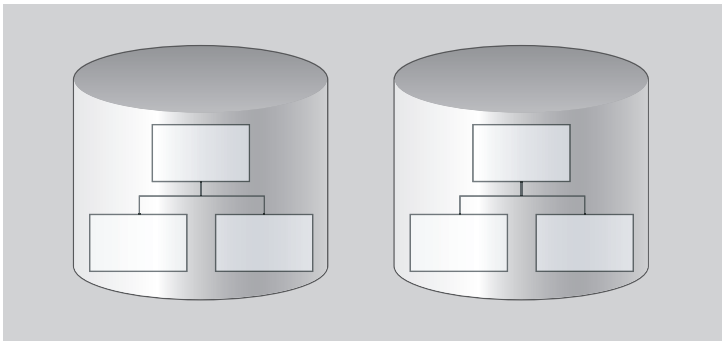

Figure 4. Creating Partitions

Symbolic Pointers

PDF requires symbolic pointing logical parent pointers, and logical children must be physically paired. Symbolic logical parent pointers and physically paired logical children are required because segment direct pointers cannot point between partitions. Logical parent, logical child, and logical twin pointers currently do not include the partition number and so they cannot be supported with any type of partitioned database.

Instead of the RBA, symbolic logical parent pointers use the concatenated key of the logical parent. Physically paired logical children store the logical child in the same partition as the physical parent, so logical child and logical twin pointers are not required.

PDF Control Statements

PDF uses the control statements described in the following sections to identify partitions, partition selection routines, and data set groups. The simplicity of converting to PDF is reflected in the easy-to-learn syntax of the control statements. These three statements are the only modifications you need to add to an IMS DBD to create a PDF database. In many cases, you need only add one or more PART statements.

PART Statements

One or more PART statements replace the first DATASET statement in the DBD of the database to be partitioned. Each PART statement identifies partitioned data set, DD name and other attributes of the partition. Up to 127 PART statements can be included in a DBD. Just as a DEDB has an AREA statement for each area, there is a PART statement for each partition.

The PART statement supports several parameters from the DATASET statement such as: DD1, SIZE, FRSPC, SCAN, and SEARCHA. It also includes two additional parameters:

- ROOT=nnn is used to define the number of blocks in the HDAM root addressable area for each partition.
- HIGHKEY=key is used to define the high key for each HIDAM partition, except for the last partition.

Note: You can use Eclipse iSurvey to produce a high-key report for those databases you plan to partition with PDF. Eclipse iSurvey is part of the NEON Eclipse Reorganization Utilities, Partitioned Support Edition. For more information about the utilities, see page 4.

PARTDBD Statements

The partition selection routine is specified by using a PARTDBD statement. The partition selection routine provides a way to specify how specific records are loaded into the partitioned database. The PARTDBD statement defines the partition selection routine to be used when the database is partitioned. You can also use a partition selection routine for secondary indexes.

DATASET Statement

The DATASET statements identify data set groups in a partition. You can include up to nine DATASET statements for each PART statement, for a total of 10 data set groups per partition.

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 Backup and Recovery Utilities

The Eclipse Backup and Recovery Utilities provide a solution for all types of IMS database recovery: point-in-time, full database recovery, and disaster recovery.

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.

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.