

IMS



Administration Guide: Database Manager

Version 9

IMS



Administration Guide: Database Manager

Version 9

Note

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 555.

Second Edition (December 2005) (Softcopy only)

This edition replaces or makes obsolete the previous edition, SC18-7806-00. This edition is available in softcopy format only. The technical changes for this version are summarized under "Summary of Changes" on page xxix.

© **Copyright International Business Machines Corporation 1974, 2005. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

| | |
|---|--------|
| Figures | xv |
| Tables | xxiii |
| About This Book | xxv |
| Prerequisite Knowledge | xxv |
| IBM Product Names Used in This Information | xxv |
| How to Read Syntax Diagrams | xxvi |
| How to Send Your Comments | xxviii |
| Summary of Changes | xxix |
| Changes to the Current Edition of This Book for IMS Version 9 | xxix |
| Changes to This Book for IMS Version 9 | xxix |
| Library Changes for IMS Version 9 | xxix |
| New and Revised Titles | xxix |
| Organizational Changes | xxx |
| Terminology Changes | xxx |
| Accessibility Enhancements | xxx |

Part 1. General Information on IMS Database Administration 1

| | |
|---|----|
| Chapter 1. Introduction to IMS Databases | 3 |
| Database Administration Overview | 3 |
| DL/I | 3 |
| CICS | 3 |
| DBCTL and DCCTL | 4 |
| Open Database Access (ODBA) | 4 |
| Database Administration Tasks | 4 |
| Concepts and Terminology | 6 |
| How Data Is Stored in a Database | 6 |
| The Hierarchy | 8 |
| The Database | 11 |
| The Database Record | 12 |
| The Segment | 14 |
| Optional Functions | 17 |
| How to Define Your Database to IMS | 18 |
| How Application Programs View the Database | 18 |
| Chapter 2. Standards and Procedures | 19 |
| Establishing Standards and Procedures | 19 |
| Naming Conventions | 21 |
| General Rules for Establishing Naming Conventions | 21 |
| HALDB Naming Conventions | 22 |
| Chapter 3. Review Process | 25 |
| The Design Review | 25 |
| Role of the Database Administrator | 25 |
| General Information about Reviews | 25 |
| Design Review 1 | 26 |
| Design Review 2 | 26 |
| Design Review 3 | 27 |
| Design Review 4 | 27 |
| Code Inspection 1 | 28 |

| | |
|---|-----------|
| Who Attends Code Inspection 1. | 28 |
| Code Inspection 2. | 28 |
| Security Inspection | 29 |
| Post-Implementation Review | 29 |
| Chapter 4. Security | 31 |
| Restricting the Scope of Data Access | 31 |
| Restricting Processing Authority. | 31 |
| Restricting Access by Non-IMS Programs | 33 |
| Protecting Data with VSAM Passwords | 33 |
| Encrypting Your Database. | 33 |
| Using the Dictionary to Help Establish Security | 33 |

Part 2. Administering IMS Databases 35

| | |
|--|-----------|
| Chapter 5. Analyzing Data Requirements | 49 |
| Local View | 49 |
| Local View 1. Current Roster. | 50 |
| Local View 2. Schedule of Classes | 51 |
| Local View 3. Instructor Skills Report. | 52 |
| Local View 4. Instructor Schedules | 53 |
| Designing a Conceptual Data Structure | 53 |
| Implementing the Structure with DL/I | 55 |
| Assigning Data Elements to Segments | 55 |
| Resolving Data Conflicts | 56 |
| Chapter 6. Choosing Full-Function Database Types | 59 |
| Sequential Storage Method | 60 |
| Direct Storage Method | 60 |
| Databases Supported with DBCTL. | 60 |
| Databases Supported with DCCTL | 61 |
| Performance Considerations Overview | 61 |
| HSAM Databases | 64 |
| When to Use HSAM | 65 |
| How an HSAM Record Is Stored | 65 |
| DL/I Calls against an HSAM Database | 67 |
| HISAM Databases | 69 |
| When to Use HISAM. | 69 |
| How a HISAM Record is Stored | 69 |
| Accessing Segments. | 72 |
| Inserting Root Segments Using VSAM | 72 |
| Inserting Dependent Segments | 74 |
| Deleting Segments | 76 |
| Replacing Segments. | 78 |
| Criteria for Selecting HISAM | 78 |
| SHSAM, SHISAM and GSAM Databases | 78 |
| Situation 1 - Converting from a non-database system to IMS | 78 |
| Situation 2 - Passing data | 79 |
| SHSAM Databases | 79 |
| SHISAM Databases | 79 |
| SHISAM IMS Symbolic Checkpoint Call. | 80 |
| GSAM Databases. | 80 |
| GSAM IMS Symbolic Checkpoint Call | 80 |
| HDAM, PHDAM, HIDAM, and PHIDAM Databases. | 82 |
| Maximum Sizes of HD Databases | 83 |
| DL/I Calls Issuable Against HD Databases. | 84 |

| | |
|---|------------|
| When to Use HDAM and PHDAM | 84 |
| When to Use HIDAM and PHIDAM | 85 |
| What You Need to Know About HD Databases | 85 |
| General Format of HD Databases and Use of Special Fields | 95 |
| How HDAM and PHDAM Records Are Stored | 98 |
| When Not Enough Root Storage Room Exists | 100 |
| How HIDAM and PHIDAM Records Are Stored | 100 |
| Accessing Segments | 103 |
| Inserting Root Segments | 104 |
| Inserting Dependent Segments | 106 |
| Deleting Segments | 107 |
| Replacing Segments | 107 |
| How the HD Space Search Algorithm Works | 107 |
| Locking Protocols | 109 |
| Managing I/O Errors | 111 |
| Chapter 7. Choosing Fast Path Database Types | 113 |
| Data Entry Databases | 113 |
| DEDB Functions | 114 |
| DEDB Areas | 114 |
| Fixed- and Variable-Length Segments in DEDBs | 120 |
| Parts of a DEDB Area | 121 |
| Root Segment Storage | 126 |
| Direct Dependent Segment Storage | 126 |
| Sequential Dependent Segment Storage | 127 |
| Enqueue Level of Segment CIs | 127 |
| DEDB Space Search Algorithm | 129 |
| DEDB Insert Algorithm | 129 |
| DEDB Free Space Algorithm | 130 |
| Managing Unusable Space with IMS Tools | 131 |
| DL/I Calls against a DEDB | 131 |
| Mixed Mode Processing | 131 |
| Main Storage Databases (MSDBs) | 132 |
| When to Use an MSDB | 133 |
| MSDBs Storage | 133 |
| MSDB Record Storage | 134 |
| Saving MSDBs for Restart | 135 |
| DL/I Calls against an MSDB | 135 |
| Rules for Using an SSA | 135 |
| Insertion and Deletion of Segments | 136 |
| Combination of Binary and Direct Access Methods | 136 |
| Position in an MSDB | 137 |
| The Field Call | 138 |
| Call Sequence Results | 138 |
| Fast Path Virtual Storage Option | 139 |
| Restrictions Using VSO DEDB Areas | 139 |
| Defining a VSO DEDB Area | 140 |
| Sharing of VSO DEDB Areas | 142 |
| Defining a VSO Cache Structure Name | 144 |
| Acquiring and Accessing Data Spaces for VSO DEDB Areas | 147 |
| Resource Control and Locking | 148 |
| Preopen Areas and VSO Areas in a Data Sharing Environment | 149 |
| Input/Output Processing With VSO | 150 |
| Checkpoint Processing | 152 |
| VSO Options Across IMS Restart | 152 |
| Emergency Restart Processing | 152 |

| | |
|--|------------|
| VSO Options with XRF | 153 |
| Fast Path Synchronization Points. | 153 |
| Phase 1 - Build Log Record. | 153 |
| Phase 2 - Write Record to System Log | 153 |
| Managing I/O Errors and Long Wait Times | 154 |
| Registering Fast Path Databases in DBRC | 155 |
| Chapter 8. Choosing Optional Database Functions | 157 |
| Logical Relationships | 157 |
| Logical Relationship Types | 158 |
| Logical Relationship Pointer Types | 162 |
| Paths in Logical Relationships | 168 |
| The Logical Child Segment | 169 |
| Segment Prefix Information for Logical Relationships | 170 |
| Intersection Data. | 170 |
| Recursive Structures: Same Database Logical Relationships | 172 |
| Defining Sequence Fields for Logical Relationships | 176 |
| Control Blocks for Logical Relationships | 177 |
| Specifying Logical Relationships in the Physical DBD | 178 |
| Specifying Logical Relationships in the Logical DBD. | 182 |
| Choosing Replace, Insert, and Delete Rules for Logical Relationships | 187 |
| Performance Considerations for Logical Relationships | 189 |
| Secondary Indexes | 192 |
| Why Secondary Indexes? | 192 |
| Characteristics of Secondary Indexes | 194 |
| Segments Used for Secondary Indexes | 194 |
| How the Hierarchy Is Restructured | 197 |
| How a Secondary Index Is Stored | 198 |
| Format and Use of Fields in a Pointer Segment | 199 |
| Making Keys Unique Using System Related Fields | 202 |
| Suppressing Index Entries: Sparse Indexing. | 204 |
| How the Secondary Index Is Maintained | 205 |
| Processing a Secondary Index as a Separate Database | 206 |
| Sharing Secondary Index Databases | 207 |
| Using the INDICES= Parameter | 207 |
| Using Secondary Indexes with Logical Relationships | 209 |
| Using Secondary Indexes with Variable-Length Segments | 210 |
| Considerations When Using Secondary Indexing | 210 |
| How to Specify Use of Secondary Indexing in the DBD. | 211 |
| Choosing Secondary Indexes Versus Logical Relationships | 214 |
| When to Use a Secondary Index | 214 |
| When to Use a Logical Relationship | 214 |
| Variable-Length Segments | 215 |
| How to Specify Variable-Length Segments | 215 |
| How Variable-Length Segments Are Stored and Processed | 216 |
| When to Use Variable-Length Segments | 217 |
| What Application Programmers Need to Know about Variable-Length Segments | 217 |
| Adding or Converting to Variable-Length Segments | 218 |
| Segment Edit/Compression Exit Routine | 218 |
| Things to Consider Before Using the Segment Edit/Compression Exit Routine | 219 |
| How to Specify the Segment Edit/Compression Exit Routine. | 220 |
| Converting to the Segment Edit/Compression Exit Routine | 220 |
| Data Capture Exit Routines | 221 |
| DBD Parameters for Data Capture Exit Routines | 221 |

| | | |
|---|--|-----|
| | Call Sequence of Data Capture Exit Routines | 223 |
| | Data Passed To and Captured By the Data Capture Exit Routine | 224 |
| | Data Capture Call Functions | 224 |
| | Cascade Delete When Crossing Logical Relationships | 225 |
| | Data Capture Exit Routines and Logically Related Databases | 225 |
| | Converting to Data Capture Exit Routines | 225 |
| | Field-Level Sensitivity | 226 |
| | Using Field-Level Sensitivity as a Mapping Interface. | 226 |
| | Using Field-Level Sensitivity with Variable-Length Segments. | 226 |
| | How to Specify Use of Field-Level Sensitivity in the DBD and PSB | 227 |
| | Retrieving Segments Using Field-Level Sensitivity | 228 |
| | Replacing Segments Using Field-Level Sensitivity | 228 |
| | Inserting Segments Using Field-Level Sensitivity | 229 |
| | Using Field-Level Sensitivity When Fields Overlap | 230 |
| | Using Field-Level Sensitivity When Path Calls Are Issued. | 230 |
| | Using Field-Level Sensitivity with Logical Relationships | 230 |
| | Using Field-Level Sensitivity with Variable-Length Segments. | 231 |
| | General Considerations for Using Field-Level Sensitivity | 235 |
| | Multiple Data Set Groups | 235 |
| | Why Use Multiple Data Set Groups? | 236 |
| | HD Databases Using Multiple Data Set Groups | 237 |
| I | VSAM KSDS CI Reclaim For Full-Function Databases | 241 |
| | HALDB Single Partition Processing | 242 |
| | Logical Relationships in Single Partition Processing | 242 |
| I | Secondary Indexes in Single Partition Processing | 242 |
| | Partition Selection | 243 |
| I | Integrated HALDB Online Reorganization Function | 244 |
| I | Storing XML Data in IMS Databases | 244 |
| | Chapter 9. Designing Full-Function Databases. | 245 |
| | Specifying Free Space (HDAM, PHDAM, HIDAM, and PHIDAM Only) | 245 |
| | Estimating the Size of the Root Addressable Area (HDAM or PHDAM Only) | 246 |
| | Determining Which Randomizing Module to Use (HDAM and PHDAM Only) | 247 |
| | Write Your Own Randomizing Module | 247 |
| | Assess the Effectiveness of the Randomizing Module | 247 |
| | Choosing HDAM or PHDAM Options | 248 |
| | Minimizing I/O Operations | 248 |
| | Maximizing Packing Density | 248 |
| | Choosing a Logical Record Length for a HISAM Database | 249 |
| | Logical Record Length Considerations | 249 |
| | Rules to Observe | 251 |
| | Calculating How Many Logical Records Are Needed to Hold a Database Record | 252 |
| | Specifying Logical Record Length | 252 |
| | Choosing a Logical Record Length for HD Databases | 252 |
| | Determining the Size of CIs and Blocks | 252 |
| | Buffering Options | 253 |
| | Multiple Buffers in Virtual Storage | 253 |
| | "Use" Chain | 253 |
| | The Buffer Handler | 253 |
| | Background Write Option. | 254 |
| | Shared Resource Pools | 254 |
| | Using Separate Subpools | 254 |
| | Hiperspace Buffering | 254 |
| | Buffer Size | 254 |
| | Buffer Numbers | 255 |

| | |
|--|------------|
| VSAM Buffer Sizes | 255 |
| OSAM Buffer Sizes | 256 |
| Specifying Buffers | 256 |
| OSAM Sequential Buffering | 257 |
| Sequential Buffering Introduction | 257 |
| Benefits of Sequential Buffering | 258 |
| Flexibility of SB Use | 259 |
| How SB Buffers Data | 259 |
| Virtual Storage Considerations for SB | 260 |
| How to Request the Use of SB | 261 |
| VSAM Options | 264 |
| Optional Functions Specified in the OPTIONS Control Statement | 264 |
| Optional Functions Specified in the POOLID, DBD, and VSRBF Control Statements | 266 |
| Optional Functions Specified in the Access Method Services DEFINE CLUSTER Command | 267 |
| OSAM Options | 269 |
| Dump Option (DUMP Parameter). | 269 |
| Deciding Which FIELD Statements to Code in the DBD | 269 |
| Planning for Maintenance | 269 |
| Chapter 10. Designing Fast Path Databases. | 271 |
| Designing a Data Entry Database (DEDB) | 271 |
| DEDB Design Guidelines. | 271 |
| DEDB Area Design Guidelines. | 272 |
| Determining the Size of the CI. | 273 |
| Determining the Size of the UOW | 274 |
| SDEP CI Preallocation and Reporting | 274 |
| Processing Option P (PROCOPT=P) | 275 |
| DEDB Randomizing Routine Design | 275 |
| Multiple Copies of an Area Data Set | 276 |
| Record Deactivation | 277 |
| Physical Child Last Pointers | 277 |
| Subset Pointers | 277 |
| Designing a Main Storage Database (MSDB) | 277 |
| Calculating Virtual Storage Requirements for an MSDB | 278 |
| Understanding Resource Allocation, a Key to Performance | 279 |
| Designing to Minimize Resource Contention. | 280 |
| Choosing MSDBs to Load and Page-Fix | 281 |
| Auxiliary Storage Requirements for an MSDB | 283 |
| High-Speed Sequential Processing (HSSP) | 283 |
| Why HSSP? | 284 |
| Limitations and Restrictions When Using HSSP | 284 |
| Using HSSP | 285 |
| HSSP Processing Option H (PROCOPT=H). | 285 |
| Image-Copy Option. | 285 |
| UOW Locking | 286 |
| Private Buffer Pools | 286 |
| Designing a DEDB or MSDB Buffer Pool | 286 |
| Buffer Requirements | 287 |
| Normal Buffer Allocation (NBA) | 287 |
| Overflow Buffer Allocation (OBA) | 287 |
| Fast Path Buffer Allocation Algorithm | 287 |
| System Buffer Allocation (DBFX) | 288 |
| Determining the Fast Path Buffer Pool Size | 288 |
| Fast Path Buffer Performance Considerations | 288 |

| | |
|--|------------|
| The NBA Limit and Sync Point | 289 |
| The DBFX Value and the Low Activity Environment | 289 |
| Designing a DEDB Buffer Pool in the DBCTL Environment | 290 |
| Buffer Requirements in a DBCTL Environment | 290 |
| Normal Buffer Allocation for BMPs | 290 |
| Normal Buffer Allocation for CCTL Regions and Threads | 290 |
| Overflow Buffer Allocation for BMPs | 291 |
| Overflow Buffer Allocation for CCTL Threads | 291 |
| Fast Path Buffer Allocation Algorithm for BMPs | 291 |
| Fast Path Buffer Allocation Algorithm for CCTL Threads | 292 |
| System Buffer Allocation (SBA) | 292 |
| Determining the Size of the Fast Path Buffer Pool for DBCTL | 292 |
| Fast Path Buffer Performance Considerations for DBCTL | 293 |
| The NBA/FPB Limit and Sync Point in a DBCTL Environment | 293 |
| Low Activity and the DBFX Value in a DBCTL Environment | 293 |
| A Note on Fast Path Buffer Allocation in IMS Regions | 294 |
| Chapter 11. Implementing Database Design | 295 |
| Coding Database Descriptions as Input for the DBDGEN Utility | 295 |
| The DBD Statement | 296 |
| The DATASET Statement | 296 |
| The SEGM Statement | 297 |
| The FIELD Statement | 297 |
| The LCHILD Statement | 297 |
| The XDFLD Statement | 298 |
| The DBDGEN and END Statements | 298 |
| Implementing HALDB Design | 298 |
| Creating HALDBs with the HALDB Partition Definition Utility | 298 |
| Allocating an ILDS | 304 |
| Coding Program Specification Blocks as Input to the PSBGEN Utility | 305 |
| The Alternate PCB Statement | 306 |
| The Database PCB Statement | 306 |
| The SENSEG Statement | 307 |
| The SENFLD Statement | 307 |
| The PSBGEN Statement | 307 |
| The END Statement | 307 |
| Building the Application Control Blocks (ACBGEN) | 308 |
| Defining Generated Program Specification Blocks for SQL Applications | 309 |
| Chapter 12. Developing Test Databases | 311 |
| Test Requirements | 311 |
| What Kind of Database? | 312 |
| What Kind of Sample Data? | 312 |
| What Kind of Application Program? | 312 |
| Designing, Creating, and Loading a Test Database | 312 |
| Using Testing Standards | 312 |
| Using IBM Programs to Develop a Test Database | 313 |
| Chapter 13. Loading Databases | 315 |
| Estimating the Minimum Size of the Database | 315 |
| Step 1. Calculate the Size of an Average Database Record | 315 |
| Step 2. Determine Overhead Needed for CI Resources | 317 |
| Step 3. Determine the Number of CIs or Blocks Needed | 318 |
| Step 4. Determine the Number of Blocks or CIs Needed for Free Space | 321 |
| Step 5. Determine the Amount of Space Needed for Bit Maps | 321 |
| Allocating Data Sets | 322 |

| | |
|--|------------|
| Allocating OSAM Data Sets | 322 |
| Example of Allocating Multi-Volume OSAM Data Sets on Non-SMS Managed DASD | 323 |
| Example of Allocating Multi-Volume OSAM Data Sets on SMS-Managed DASD | 324 |
| Cautions When Allocating OSAM Data Sets | 324 |
| Writing a Load Program | 325 |
| The Load Process | 325 |
| Status Codes for Load Programs | 326 |
| Using SSAs in a Load Program | 327 |
| Loading a Sequence of Segments with the D Command Code | 327 |
| Loading a HISAM Database | 336 |
| Loading a SHISAM Database | 337 |
| Loading a GSAM Database | 337 |
| Loading an HDAM or a PHDAM Database | 337 |
| Loading a HIDAM or a PHIDAM Database | 337 |
| Loading a Database with Logical Relationships or Secondary Indexes | 337 |
| Loading Fast Path Databases | 337 |
| Loading an MSDB | 337 |
| Loading a DEDB | 338 |
| Loading Sequential Dependent Segments | 339 |
| Chapter 14. Monitoring Databases | 341 |
| IMS Monitor | 341 |
| Monitoring Fast Path Systems | 343 |
| Fast Path Log Analysis Utility | 343 |
| Interpreting Fast Path Analysis Reports | 345 |
| Chapter 15. Tuning Databases | 347 |
| Reorganizing the Database | 347 |
| When You Should Reorganize | 348 |
| Reorganizing Databases Offline | 348 |
| Protecting Your Database During an Offline Reorganization | 348 |
| Offline Reorganization Utilities | 349 |
| Procedures for Offline Database Reorganizations | 363 |
| Reorganizing HALDBs | 364 |
| HALDB Offline Reorganization | 365 |
| HALDB Online Reorganization | 370 |
| The HALDB Self-Healing Pointer Process | 388 |
| Changing DL/I Access Methods | 394 |
| Changing the DL/I Access Method From HISAM to HIDAM | 395 |
| Changing the DL/I Access Method From HISAM to HDAM | 395 |
| Changing the DL/I Access Method From HIDAM to HISAM | 397 |
| Changing the DL/I Access Method From HIDAM to HDAM | 397 |
| Changing the DL/I Access Method From HDAM to HISAM | 398 |
| Changing the DL/I Access Method From HDAM to HIDAM | 399 |
| Changing the DL/I Access Method From HDAM to PHDAM and HIDAM to PHIDAM | 401 |
| Changing HALDB Partition Definitions | 404 |
| Procedure for Changing to DEDBs | 407 |
| Changing the Hierarchic Structure | 408 |
| Changing the Sequence of Segment Types | 408 |
| Combining Segments | 408 |
| Procedure for Changing the Hierarchic Structure | 409 |
| Changing Direct-Access Storage Devices | 409 |
| Tuning OSAM Sequential Buffering | 410 |

| | |
|---|------------|
| Well-Organized Database | 410 |
| Badly-Organized Database | 410 |
| Ensuring a Well-Organized Database | 410 |
| Adjusting HDAM and PHDAM Options | 411 |
| Adjusting Buffers | 411 |
| VSAM Buffers | 411 |
| OSAM Buffers | 413 |
| Procedure for Adjusting VSAM and OSAM Database Buffers | 413 |
| OSAM Sequential Buffering | 414 |
| Procedure for Adjusting Sequential Buffers | 414 |
| Adjusting VSAM Options | 414 |
| Procedure for Adjusting VSAM Options Specified in the OPTIONS Control Statement | 415 |
| Procedures for Adjusting VSAM Options Specified in the Access Method Service DEFINE CLUSTER Command | 415 |
| Adjusting OSAM Options | 416 |
| Changing the Amount of Space Allocated | 416 |
| Changing Operating System Access Methods | 417 |
| Changing the Number of Data Set Groups | 418 |
| Tuning Fast Path Systems | 421 |
| Transaction Volume to a Particular Fast Path Application Program | 422 |
| DEDB Structure Considerations | 422 |
| Usage of Buffers from a Buffer Pool | 423 |
| Contention for DEDB Control Interval (CI) Resources | 424 |
| Exhaustion of DEDB DASD Space | 425 |
| Utilization of Available Real Storage | 425 |
| Synchronization Point Processing and Physical Logging | 425 |
| Contention for Output Threads | 425 |
| Overhead Resulting from Reprocessing | 426 |
| Dispatching Priority of Processor-Dominant and I/O-Dominant Tasks | 426 |
| DASD Contention Due to I/O on DEDBs | 426 |
| Resource Locking Considerations with Block Level Sharing | 426 |
| Resource Name Hash Routine | 427 |
| Chapter 16. Modifying Databases | 429 |
| Adding Segment Types | 430 |
| Unloading and Reloading Using the Reorganization Utilities | 430 |
| Without Unloading or Reloading | 431 |
| Using Your Own Unload and Reload Program | 431 |
| Deleting Segment Types | 431 |
| Moving Segment Types | 432 |
| Changing Segment Size | 432 |
| Changing Data in a Segment (Except for Data at the End of a Segment) | 433 |
| Changing the Position of Data in a Segment | 433 |
| Adding Logical Relationships | 433 |
| Example 1. DBX Exists, DBY Is to Be Added | 434 |
| Example 2. DBX and DBY Exist, DBZ Is to Be Added | 435 |
| Example 3. DBX and DBY Exist, DBZ Is to Be Added | 436 |
| Example 4. DBX and DBY Exist, DBZ Is to Be Added | 436 |
| Example 5. DBX Exists, DBY Is to Be Added | 437 |
| Example 6. DBX and DBY Exist, DBZ Is to Be Added | 438 |
| Example 7. DBX and DBY Exist, DBZ Is to Be Added | 440 |
| Example 8. DBX and DBY Exist, DBZ Is to Be Added | 442 |
| Example 9. DBY Exists, DBZ Is to Be Added | 442 |
| Example 10. DBY Exists, DBZ Is to Be Added | 443 |
| Example 11. DBX and DBY Exist, DBZ Is to Be Added | 443 |

| | |
|--|-----|
| Example 12. DBX and DBY Exist, DBZ Is to Be Added. | 444 |
| Example 13. DBX and DBY Exist, Segment Y and DBZ Are to Be Added | 444 |
| Steps in Reorganizing a Database to Add a Logical Relationship | 445 |
| Some Restrictions on Modifying Existing Logical Relationships | 449 |
| Summary on Use of Utilities When Adding Logical Relationships | 450 |
| Adding a Secondary Index | 451 |
| Adding or Converting to Variable-Length Segments | 451 |
| Method 1. Converting Segments or a Database | 451 |
| Method 2. Converting Segments or a Database | 452 |
| Converting to the Segment Edit/Compression Exit Routine | 452 |
| Converting Databases for Data Capture Exit Routines and Asynchronous Data Capture | 453 |
| Converting a Logical Parent Concatenated Key from Virtual to Physical or Physical to Virtual | 454 |
| Using the Online Change Function | 454 |
| Online Change, DEDBs, and Availability of IFP and MPP Regions | 455 |
| Online Change and DEDB Randomizer and Exit Routines | 457 |
| Making Online Changes at the DEDB and Area Level | 461 |
| Extending DEDB Independent Overflow Online | 464 |

| | |
|-------------------------------------|------------|
| Part 3. Appendixes | 467 |
|-------------------------------------|------------|

| | |
|---|------------|
| Appendix A. Meaning of Bits in the Delete Byte | 469 |
| Bits in the Delete Byte. | 469 |
| Bits in the Prefix Descriptor Byte | 469 |

| | |
|--|------------|
| Appendix B. Insert, Delete, and Replace Rules for Logical Relationships | 471 |
| Specifying Rules in the Physical DBD | 471 |
| Insert Rules | 472 |
| The Logical Child Insert Call | 472 |
| Status Codes | 473 |
| Insert Rules Summary. | 475 |
| Replace Rules | 475 |
| The Replace Call | 475 |
| Replace Rule Status Codes. | 475 |
| Replace Rules Summary. | 478 |
| Using the DLET Call | 481 |
| Physical and Logical Deletion | 481 |
| Deleting Concatenated Segments | 481 |
| The Third Access Path | 482 |
| Use of the Delete Byte | 483 |
| Issuing the Delete Call | 483 |
| Status Codes | 483 |
| DASD Space Release. | 484 |
| Delete Rules | 484 |

| | |
|---|------------|
| Appendix C. Using OSAM as the Access Method. | 513 |
|---|------------|

| | |
|--|------------|
| Appendix D. Correcting Bad Pointers | 515 |
|--|------------|

| | |
|---|------------|
| Appendix E. HALDB Partition Definition utility | 517 |
| The Partitioned Databases Panel. | 518 |
| Accessing Help Information. | 519 |
| Exiting the Utility. | 519 |
| Displaying the ISPF Member List. | 520 |
| Opening HALDB Partitions | 521 |

| | | |
|--|--|-----|
| | Automatic Partition Definition | 527 |
| | Manual Partition Definition | 528 |
| | Defining Data Set Group Information | 533 |
| | Displaying the List of Defined Partitions | 534 |
| | The Partition List Line Commands | 538 |
| | The Partition List Action Bar | 538 |
| | Change All Partitions | 540 |
| | Opening Database Information | 542 |
| | Deleting Database Information. | 543 |
| | Exporting Database Information | 543 |
| | Importing Database Information | 544 |
| | Displaying the IMS Concatenation | 544 |
| | Selecting an IMS Configuration Using the Partition Definition Utility | 545 |
| | Using Batch to Export or Import Partition Information | 547 |
| | DSPXRUN Command Syntax | 548 |
| | DSPXRUN EXPORT Sample Output | 549 |
| | DSPXRUN IMPORT Sample Output | 549 |
| | Appendix F. Output Data Set Requirements for HALDB Online | |
| | Reorganization | 551 |
| | HALDB Online Reorganization Requirements for Existing Output Data Sets | 551 |
| | Attributes of Automatically-Created Output Data Sets | 551 |
| | Notices | 555 |
| | Programming Interface Information | 557 |
| | Trademarks. | 558 |
| | Bibliography | 559 |
| | IMS Version 9 Library | 559 |
| | Supplementary Publications. | 560 |
| | Publication Collections | 560 |
| | Accessibility Titles Cited in This Library | 560 |
| | Index | 561 |

Figures

| | |
|---|----|
| 1. Segment Types in the School Database Record | 7 |
| 2. Segment Occurrences in a School Database Record | 8 |
| 3. Hierarchic Sequence of Segment Types for School Database | 9 |
| 4. Hierarchic Sequence of Segment Occurrences for School Database | 10 |
| 5. Levels in the Database | 11 |
| 6. An Example of a Medical Database Record | 12 |
| 7. Example of Records That Can Be Stored in the School Database | 13 |
| 8. Records that Cannot be Stored in the School Database | 13 |
| 9. Format of Fixed-Length Segments | 14 |
| 10. Format of Variable-Length Segments | 14 |
| 11. Three Segment Occurrences and Three Data Elements of the STUDENT Segment | 16 |
| 12. Example of STUDENT Segments Stored in Alphabetic Order | 16 |
| 13. DBD for Payroll Database | 32 |
| 14. Payroll Database Record without a Mask | 32 |
| 15. PCB for Payroll Database | 32 |
| 16. Payroll Database Record with SALARY Segment Masked | 33 |
| 17. Current Roster Conceptual Data Structure | 51 |
| 18. Schedule of Classes Conceptual Data Structure | 52 |
| 19. Instructor Skills Report Conceptual Data Structure | 52 |
| 20. Instructor Schedules Conceptual Data Structure | 53 |
| 21. Education Data Structures | 55 |
| 22. Education Hierarchies | 57 |
| 23. Bidirectional Logical Relationships | 57 |
| 24. Example HSAM Database | 66 |
| 25. Example HSAM Database Stored in Blocks | 66 |
| 26. GU Calls against an HSAM Database | 67 |
| 27. Updating an HSAM Database | 68 |
| 28. Example HISAM Database | 70 |
| 29. Example HISAM Database in Storage | 71 |
| 30. Format of a Logical Record in a HISAM Database | 72 |
| 31. Inserting a Root Segment into a HISAM Database (Free Logical Record Exists in the CI) | 73 |
| 32. Inserting a Root Segment into a HISAM Database (No Free Logical Record Exists in the CI) | 74 |
| 33. Inserting a Dependent Segment into a HISAM Database (Space Exists in the Logical Record) | 75 |
| 34. Inserting a Dependent Segment into a HISAM Database (No Space Exists in the Logical Record) | 76 |
| 35. The Hierarchic Segment Layout on the Database | 77 |
| 36. Accessing a HISAM Segment That Hierarchically Follows Deleted Segments. | 77 |
| 37. A Logical View of an HDAM and a PHDAM Database | 82 |
| 38. A Logical View of a HIDAM and a PHIDAM | 83 |
| 39. Example Database Record | 85 |
| 40. Example Database Record for Illustrating Pointers | 86 |
| 41. Hierarchic Forward Pointers | 87 |
| 42. Hierarchic Forward and Backward Pointers | 88 |
| 43. Physical Child First Pointers | 89 |
| 44. Physical Child First and Last Pointers | 90 |
| 45. Specifying PCF and PCL Pointers | 91 |
| 46. Physical Twin Forward Pointers | 92 |
| 47. Physical Twin Forward and Backward Pointers | 92 |
| 48. Mixing Pointers | 94 |
| 49. Format of an HD Database and Special Fields in It | 95 |
| 50. Bit Map for HD Databases | 96 |
| 51. An FSEAP | 97 |
| 52. An FSE | 97 |
| 53. An HDAM or PHDAM Anchor Point Area | 98 |

| | |
|---|-----|
| 54. Two Example HD Database Records | 98 |
| 55. HDAM or PHDAM Database Records in Storage | 99 |
| 56. HIDAM or PHIDAM Database Records in Storage | 101 |
| 57. Format of an Index Segment | 102 |
| 58. HIDAM or PHIDAM Index Databases | 102 |
| 59. Specifying PTR=T or PTR=H for Root Segments in a HIDAM Database | 103 |
| 60. How Dependent Segments Are Found Using PCF and PTF Pointers | 104 |
| 61. Inserting a Root Segment into a HIDAM or PHIDAM Database | 105 |
| 62. Updating the Space Management Fields in an HDAM or PHDAM Database. | 106 |
| 63. Defining a Variable-Length Segment | 120 |
| 64. Defining a Fixed-Length Segment | 120 |
| 65. Parts of a DEDB Area in Storage | 122 |
| 66. CI Format | 123 |
| 67. Root Segment Format (with Sequential and Direct Dependent Segments with Subset Pointers) | 124 |
| 68. Sequential Dependent Segment Format | 125 |
| 69. Direct Dependent Segment Format. | 125 |
| 70. DEDB Structure Example | 128 |
| 71. Extending a UOW to Use Independent Overflow. | 130 |
| 72. MSDB Pointers | 133 |
| 73. MSDBINIT Record Format | 134 |
| 74. Sequence of the Four MSDB Organizations | 135 |
| 75. ECNT and MSDB Storage Layout | 137 |
| 76. Example of Updating a Policy with New Structures | 145 |
| 77. Defining a VSO Area Coupling Facility Structure Name in DBRC. | 146 |
| 78. Examples of Defining Private Buffer Pools | 146 |
| 79. A Simple Logical Relationship | 158 |
| 80. Unidirectional Logical Relationship | 159 |
| 81. Two Unidirectional Logical Relationships. | 160 |
| 82. Bidirectional Physically Paired Logical Relationship. | 160 |
| 83. Bidirectionally Virtually Paired Logical Relationship | 161 |
| 84. Direct Logical Parent (LP) Pointer | 163 |
| 85. Symbolic Logical Parent (LP) Pointer | 164 |
| 86. Logical Child First (LCF) Pointer (Used in Virtual Pairing Only) | 165 |
| 87. Physical Parent (PP) Pointer | 166 |
| 88. Logical Twin Forward (LTF) Pointer (Used in Virtual Pairing Only) | 167 |
| 89. Self-healing Pointers | 167 |
| 90. Defining a Physical Parent to Logical Parent Path in a Logical Database. | 168 |
| 91. Defining a Logical Parent to Physical Parent Path in a Logical Database. | 168 |
| 92. Format of a Concatenated Segment Returned to User I/O Area | 169 |
| 93. Fixed Intersection Data | 171 |
| 94. Variable Intersection Data | 172 |
| 95. Model 1 Components and Subassemblies | 174 |
| 96. Database Records for the Model 1 Bicycle | 175 |
| 97. Extra Database Records Required for the Model 2 Bicycle | 176 |
| 98. Relationship of Control Blocks When a Logical Relationship Is Used | 178 |
| 99. Layouts of Segments Used in the Examples | 179 |
| 100. Physical DBDs for Unidirectional Relationship Using Symbolic Pointing | 179 |
| 101. Logical Data Structure for a Unidirectional Relationship Using Symbolic Pointing | 182 |
| 102. Definition of Crossing a Logical Relationship | 184 |
| 103. The First Logical Relationship Crossed in a Hierarchic Path of a Logical Database | 185 |
| 104. Logical Database Hierarchy Enabled by Crossing the First Logical Relationship | 186 |
| 105. Single Concatenated Segment Type Defined Multiple Times with Different Combinations of Key and Data Sensitivity | 187 |
| 106. Example of the Replace, Insert, and Delete Rules | 188 |
| 107. Example of the Replace, Insert, and Delete Rules: Before and After | 188 |
| 108. Example of a Unidirectional Logical Relationship. | 191 |

| | |
|---|-----|
| 109. Example of a Logical Structure | 191 |
| 110. Database Record in Educational Database | 193 |
| 111. Example of a Database Record Unique Key Field | 193 |
| 112. Segments Used for Secondary Indexes | 194 |
| 113. Format of Pointer Segments Contained in the Secondary Index Database | 195 |
| 114. Education Database Record | 196 |
| 115. How a Segment Is Accessed Using a Secondary Index | 196 |
| 116. Call an Application Issues When a Secondary Index Is Used | 196 |
| 117. Physical Database Structure with Target Segment G | 197 |
| 118. Secondary Index Structure Indexed in Secondary Index on Segment G | 198 |
| 119. Examples of Source Segments for Each Student | 199 |
| 120. Example of a Logical Record Containing a Pointer Segment | 199 |
| I 121. Secondary Index Entry for HALDB | 199 |
| 122. Examples of Several Source Segments for Each Student | 200 |
| 123. Database Record Showing the Source and Target for Secondary Indexes | 203 |
| 124. Concatenated Key of the STUDENT Segment | 203 |
| 125. Databases for First Example of the INDICES= Parameter | 208 |
| 126. PCB for the First Example of the INDICES= Parameter | 208 |
| 127. Application Program Call Issued for the First Example of the INDICES= Parameter | 208 |
| 128. Databases for Second Example of the INDICES= Parameter | 209 |
| 129. PCB for the Second Example of the INDICES= Parameter | 209 |
| 130. Application Program Call Issued for the Second Example of the INDICES= Parameter. | 210 |
| 131. Databases for Secondary Indexing Example | 212 |
| 132. EDUC DBD for Secondary Indexing | 213 |
| 133. SINDX DBD for Secondary Indexing | 213 |
| 134. Fields in the CUSTOMER Segment | 214 |
| 135. Assembly and Parts as Examples to Demonstrate Segments Logical Relationship | 214 |
| 136. Example of a Segment That Appears to Have Two Parents | 215 |
| 137. How Variable-Length Segments Are Specified. | 215 |
| 138. Format of HISAM Variable-Length Segments | 216 |
| 139. Format of HDAM, PHDAM, HIDAM or PHIDAM Variable-Length Segments | 216 |
| 140. DBD and PSB Coding for Field-Level Sensitivity | 227 |
| 141. DBD Example for Field-Level Sensitivity | 227 |
| 142. PSB Example for Field-Level Sensitivity | 227 |
| 143. Example of a Retrieve Call. | 228 |
| 144. Example of a REPL Call | 229 |
| 145. Example of an ISRT Call | 230 |
| 146. Example of a Missing Field on a Retrieve Call | 231 |
| 147. DBD Example for Field-Level Sensitivity with Variable-Length Segments | 232 |
| 148. PSB Example for Field-Level Sensitivity with Variable-Length Segments | 232 |
| 149. First Example of a Missing Field on a Replace Call. | 232 |
| 150. Second Example of a Missing Field on a Replace Call | 233 |
| 151. Example of a Missing Field on an Insert Call | 234 |
| 152. Example of a Partially Present Field on a Retrieval Call | 234 |
| 153. Example of a Partially Present Field on a REPL Call | 235 |
| 154. Hierarchy of Applications That Need to Access INSTR and LOC Segments | 236 |
| 155. Database Record Split into Two Database Groups | 237 |
| 156. Example of How to Divide an HD Database Record | 238 |
| 157. Connecting Segments in Multiple Data Set Groups Using Physical Child First Pointers. | 238 |
| 158. HD Database Record in Storage When Multiple Data Set Groups Are Used. | 239 |
| 159. First Example of Data Set Groups | 240 |
| 160. HDAM DBD for First Example of Data Set Groups | 240 |
| 161. PHDAM DBD for First Example of Data Set Groups | 240 |
| 162. Second Example of Data Set Groups | 241 |
| 163. HDAM DBD for Second Example of Data Set Groups | 241 |
| 164. PHDAM DBD for Second Example of Data Set Groups | 241 |

| | |
|---|-----|
| 165. Specifying the RMNAME keyword | 248 |
| 166. Database Record for Logical Record Examples | 250 |
| 167. Short Logical Records | 250 |
| 168. Long Logical Records | 250 |
| 169. Database Record for Logical Records Example | 251 |
| 170. Logical Records Example with Two Read Operations | 251 |
| 171. Levels in a VSAM Index. | 268 |
| 172. First Example MSDB Record Held in Exclusive Mode | 281 |
| 173. Second Example MSDB Record Held in Exclusive Mode. | 281 |
| 174. The DBD Generation Process | 296 |
| 175. Structure of DBD Generation Input. | 296 |
| 176. Example of a Date Field within a Segment Defined as Three 2-Byte Fields and One 6-Byte Field | 297 |
| 177. Partition Default Information | 300 |
| I 178. Change Partition Panel | 302 |
| 179. Sample Command to Define an ILDS | 305 |
| 180. The PSB Generation Process. | 306 |
| 181. Structure of PSB Generation Input | 306 |
| 182. Example of a SENSEG Relationship | 307 |
| 183. The ACB Generation Process. | 308 |
| 184. Segment Sizes and Average Segment Occurrences | 317 |
| 185. JCL allocating a multi-volume OSAM data set on non-SMS managed DASD | 323 |
| I 186. Example JCL for Allocating a Multi-Volume OSAM Data Set on SMS-Managed DASD | 324 |
| I 187. Example of ALLOCATE Command for Allocating a Multi-Volume OSAM Data Set on SMS-Managed DASD | 324 |
| 188. The Load Process | 328 |
| 189. Loading a Database Using Existing Files | 329 |
| 190. Basic Initial Load Program Logic | 331 |
| 191. Sample Load Program | 332 |
| 192. Restartable Initial Load Program Logic | 334 |
| 193. Sample Restartable Initial Load Program | 335 |
| 194. JCL used to initially load a database | 336 |
| 195. IMS Monitor Works | 342 |
| 196. Fast Path Transaction Event Timings | 344 |
| 197. Steps in Reorganizing When Logical Relationships or Secondary Indexes Exist | 352 |
| I 198. Steps for Reorganizing HALDB Partitions When Logical Relationships or Secondary Indexes Exist | 353 |
| 199. HISAM Reorganization Unload Utility (DFSURUL0). | 353 |
| 200. HISAM Reorganization Reload Utility (DFSURRL0). | 354 |
| 201. HD Reorganization Unload Utility (DFSURGU0) | 354 |
| 202. HD Reorganization Reload Utility (DFSURGL0) | 355 |
| 203. Database Prereorganization Utility (DFSURPR0). | 356 |
| 204. Database Scan Utility (DFSURGS0) | 357 |
| 205. Database Prefix Resolution Utility (DFSURG10) | 358 |
| 206. Database Prefix Update Utility (DFSURGP0) | 359 |
| 207. HISAM Reorganization Unload and Reload Utilities Used for Create, Merge, or Replace Secondary Indexing Operations | 360 |
| 208. HISAM Reorganization Unload Utility Used for Extract Secondary Indexing Operations | 361 |
| 209. Database Surveyor Utility (DFSPRSUR). | 362 |
| 210. Partial Database Reorganization Utility (DFSPRCT1) | 363 |
| I 211. Offline Reorganization of a HALDB database | 366 |
| I 212. Example: The HD Reorganization Unload Utility Control Statement to Unload One Partition | 367 |
| I 213. Example: The HD Reorganization Unload Utility Control Statement to Unload Multiple Partitions | 367 |
| I 214. Example: Sample JCL to Unload a HALDB Partition | 368 |
| I 215. Example: IEC161I message during reload | 368 |
| I 216. Example: JCL to Reload a HALDB Partition | 369 |

| | |
|---|-----|
| 217. Example RECON Listing: DB Record for a HALDB in Cursor-Active Status | 372 |
| 218. The Relationship between Input Data Sets and Output Data Sets during the Online Reorganization of a HALDB Partition | 373 |
| 219. Normal Processing Steps of HALDB Online Reorganization. | 375 |
| 220. Processing Steps for an Interrupted Online Reorganization of a HALDB Partition. | 382 |
| 221. HALDB Pointer Before a Reorganization. | 389 |
| 222. HALDB Pointer After a Reorganization | 391 |
| 223. HALDB Pointer After the Self-Healing Process | 392 |
| 224. HDAM and HIDAM Databases Before and After Changing to PHDAM and PHIDAM. | 401 |
| 225. Utility Sequence of Execution When Making Database Changes during Reorganization | 419 |
| 226. Where Segment Types Can Be Added in a Database Record | 430 |
| 227. DBX Exists, DBY Is to Be Added | 434 |
| 228. DBX and DBY Exist, DBZ Is to Be Added | 435 |
| 229. DBX and DBY Exist, DBZ Is to Be Added | 436 |
| 230. DBX and DBY Exist, DBZ Is to Be Added | 437 |
| 231. DBX Exists and DBY Is to Be Added | 437 |
| 232. DBX and DBY Exist, DBZ Is to Be Added | 438 |
| 233. DBX and DBY Exist, DBZ Is to Be Added | 440 |
| 234. DBX and DBY Exist, DBZ Is to Be Added | 442 |
| 235. DBY Exists, DBZ Is to Be Added | 442 |
| 236. DBY Exists, DBZ Is to Be Added | 443 |
| 237. DBX and DBY Exist, DBZ Is to Be Added | 443 |
| 238. DBX and DBY Exist, DBZ Is to Be Added | 444 |
| 239. DBX and DBY Exist, Segment Y and DBZ Are to Be Added | 444 |
| 240. The Change in Pairing from Virtual to Physical | 450 |
| 241. The Position Change of a Real Logical Child from One Logically Related Database to Another | 450 |
| 242. Adding a Database Using Online Change | 461 |
| 243. Insert, Delete, and Replace Rules in the DBD. | 471 |
| 244. Physical Insert Rule Example. | 473 |
| 245. Paths for Physical Insert Rule Example | 473 |
| 246. ISRT and Status Codes for Physical Insert Rule Example | 473 |
| 247. Logical Insert Rule Example | 474 |
| 248. ISRT and Status Codes for Logical Insert Rule Example | 474 |
| 249. Virtual Insert Rule Example | 474 |
| 250. ISRT and Status Codes for Virtual Insert Rule Example | 474 |
| 251. Physical Replace Rule Example. | 476 |
| 252. Calls and Status Codes for Physical Replace Rule Example | 476 |
| 253. Logical Replace Rule Example | 477 |
| 254. Calls and Status Codes for Logical Replace Rule Example | 477 |
| 255. Virtual Replace Rule Example | 478 |
| 256. Calls and Status Codes for Virtual Replace Rule Example | 478 |
| 257. Physical Databases for Replace Rules Tables. | 479 |
| 258. Logical Views for Replace Rules Table | 479 |
| 259. Concatenated Segment Relationships. | 482 |
| 260. Third Access Path Example | 482 |
| 261. Logical Parent, Virtual Pairing—Physical Delete Rule Example | 485 |
| 262. Logical Parent, Physical Pairing—Physical Delete Rule Example: Before and After | 486 |
| 263. Logical Parent, Physical Pairing—Physical Delete Rule Example: Database Calls | 486 |
| 264. Logical Parent, Physical Pairing—Physical Delete Rule Example. | 486 |
| 265. Logical Parent, Physical Pairing—Physical Delete Rule Example: Before and After | 487 |
| 266. Logical Parent, Physical Pairing—Physical Delete Rule Example: Calls and Status Codes | 487 |
| 267. Logical Parent, Virtual Pairing—Logical Delete Rule Example | 487 |
| 268. Logical Parent, Virtual Pairing—Logical Delete Rule Example: Before and After | 488 |
| 269. Logical Parent, Virtual Pairing—Logical Delete Rule Example: Calls and Status Codes | 488 |
| 270. Logical Parent, Physical Pairing—Logical Delete Rule Example | 489 |
| 271. Logical Parent, Physical Pairing—Logical Delete Rule Example: Before and After | 489 |

| | |
|--|-----|
| 272. Logical Parent, Physical Pairing—Logical Delete Rule Example: Calls and Status Codes | 489 |
| 273. Logical Parent, Virtual Pairing—Virtual Delete Rule Example | 490 |
| 274. Logical Parent, Virtual Pairing—Virtual Delete Rule Example: Before and After | 490 |
| 275. Logical Parent, Virtual Pairing—Virtual Delete Rule Example: Calls and Status Codes | 490 |
| 276. Logical Parent, Physical Pairing—Virtual Delete Rule Example | 491 |
| 277. Logical Parent, Physical Pairing—Virtual Delete Rule Example: Before and After | 491 |
| 278. Logical Parent, Physical Pairing—Virtual Delete Rule Example: Calls and Status | 491 |
| 279. Physical Parent, Virtual Pairing—Bidirectional Virtual Example. | 492 |
| 280. Physical Parent, Virtual Pairing—Bidirectional Virtual Example: Before and After | 492 |
| 281. Deleting Last Logical Child Deletes Physical Parent | 492 |
| 282. Logical Child, Virtual Pairing—Physical Delete Rule Example | 493 |
| 283. Logical Child, Virtual Pairing—Physical Delete Rule Example: Deleting the Logical Child | 493 |
| 284. Logical Child, Virtual Pairing—Physical Delete Rule Example: Before and After | 494 |
| 285. Logical Child, Virtual Pairing—Logical Delete Rule Example | 494 |
| 286. Logical Child, Virtual Pairing—Logical Delete Rule Example: Calls and Status | 495 |
| 287. Logical Child, Virtual Pairing—Logical Delete Rule Example: Before and After | 495 |
| 288. Logical Child, Physical Pairing—Physical or Logical Delete Rule Example | 496 |
| 289. Logical Child, Physical Pairing—Physical or Logical Delete Rule Example: Calls and Status | 496 |
| 290. Logical Child, Physical Pairing—Physical or Logical Delete Rule Example: Before and After | 497 |
| 291. Logical Child, Virtual Pairing—Virtual Delete Rule Example | 497 |
| 292. Logical Child, Virtual Pairing—Virtual Delete Rule Example: Calls and Status | 498 |
| 293. Logical Child, Virtual Pairing—Virtual Delete Rule Example: Before and After | 498 |
| 294. Logical Child, Physical Pairing—Virtual Delete Rule Example | 499 |
| 295. Logical Child, Physical Pairing—Virtual Delete Rule Example: Calls and Status | 499 |
| 296. Logical Child, Physical Pairing—Virtual Delete Rule Example: Before and After | 500 |
| 297. (Part 1 of 5). Example of Deleted Segments Accessibility | 501 |
| 298. (Part 2 of 5). Example of Deleted Segments Accessibility | 502 |
| 299. (Part 3 of 5). Example of Deleted Segments Accessibility | 502 |
| 300. (Part 4 of 5). Example of Deleted Segments Accessibility: Database Calls | 503 |
| 301. (Part 5 of 5). Example of Deleted Segments Accessibility | 503 |
| 302. Example of Abnormal Termination | 504 |
| 303. Example of Violation of the Physical Delete Rule | 505 |
| 304. Example of Violation of the Physical Delete Rule: Database Calls | 505 |
| 305. Example of Treating the Physical Delete Rule as Logical. | 506 |
| 306. Example of Treating the Physical Delete Rule as Logical: Database Calls | 506 |
| 307. Insert, Delete, and Replace Rules Summary | 509 |
| 308. Partitioned Databases panel (DSPXPAA) | 518 |
| 309. Help Action Bar Choices | 519 |
| 310. Exit Confirmation Panel | 520 |
| 311. ISPF Member List Display (DSPXPAM) | 520 |
| 312. File Action Bar Choices | 521 |
| I 313. Partitioned Database Information (DSPXPOA) | 522 |
| 314. Partition Default Information (DSPXPCA) | 524 |
| 315. Automatic Definition Status | 528 |
| I 316. Change Partition (DSPXPPA) | 530 |
| I 317. Selection String Editor (DSPXPKE) | 532 |
| 318. Change Data Set Groups, Part 1 (DSPXPGA) | 533 |
| 319. Change Data Set Groups, Part 2 (DSPXPGB) | 534 |
| 320. Change a Data Set Group (DSPXPGC) | 534 |
| I 321. Database Partitions Panel, Sorted by Partition ID (DSPXPLA) | 535 |
| I 322. Database Partitions Panel, Sorted by Key (DSPXPLB) | 536 |
| 323. Database Partitions Panel, Sorted by Name (DSPXPLC). | 537 |
| 324. File Action Bar Choices | 538 |
| 325. Edit Action Bar Choices | 539 |
| 326. Searching the Partition List | 540 |
| 327. View Action Bar Choices | 540 |

| | | |
|---|---|-----|
| | 328. Change Partition Panel (DSPXPPB) | 541 |
| | 329. Change Data Set Groups, Part 1 (DSPXPGA) | 542 |
| I | 330. Partitioned Database Information (DSPXPOA) | 542 |
| | 331. Delete a Database (DSPXPDA) | 543 |
| | 332. Export a Database (DSPXPEA) | 543 |
| | 333. Import a Database (DSPXPIA) | 544 |
| | 334. The IMS Concatenation (ISRDDN) | 545 |
| | 335. User Configurations (DSPXPMB) | 546 |
| | 336. Configuration Details Panel (DSPXPMC) | 546 |
| I | 337. Sample JCL for Batch Import | 547 |

Tables

| | | |
|-----|---|-----|
| 1. | Licensed Program Full Names and Short Names. | xxv |
| 2. | Types of IMS Databases and the z/OS Access Methods They Can Use. | 11 |
| 3. | Example of Naming Conventions | 22 |
| 4. | Suffixes for DD names. | 23 |
| 5. | Combined Mappings for Local Views | 54 |
| 6. | Keys and Associated Data Elements | 55 |
| 7. | Summary of Database Characteristics and Options for Database Types | 63 |
| 8. | Comparison of SHSAM, SHISAM, and GSAM Databases | 81 |
| 9. | Maximum Sizes for HDAM, HIDAM, PHDAM, and PHIDAM Databases | 83 |
| 10. | CI Format | 124 |
| 11. | Root Segment Format | 124 |
| 12. | Sequential Dependent Segment Format | 125 |
| 13. | Direct Dependent Segment Format. | 125 |
| 14. | MSDBINIT Record Format | 134 |
| 15. | Required CFRM List Structure Storage Sizes for CFLEVEL=12 | 154 |
| 16. | Parts List for the Model 1 Bicycle Example | 173 |
| 17. | Delete Rule Restrictions for Logically Related Databases Using Data Capture Exit Routines | 225 |
| 18. | Examples of Multiple Data Set Grouping. | 237 |
| 19. | Levels of Enqueue of an MSDB Record | 279 |
| 20. | Example of MSDB Record Status: Shared (S) or Owned Exclusively (E) | 279 |
| 21. | File Names and Data Sets to Allocate. | 299 |
| 22. | Minimum and maximum number of data sets for HALDB partitions. | 303 |
| 23. | Required Fields and Pointers in a Segment's Prefix | 316 |
| 24. | Calculating the Average Database Record Size | 317 |
| 25. | VSAM Control Fields | 318 |
| 26. | Monitor Data for Fast Path Transactions. | 345 |
| 27. | IMS Versions that Can Access HALDBs that Are Capable of Being Reorganized Online | 376 |
| 28. | Data Set Name Examples for HALDB Online Reorganization | 379 |
| 29. | Mapping Startup Tasks to Commands for HALDB Online Reorganization | 379 |
| 30. | Mapping Monitoring Tasks to Commands for HALDB Online Reorganization | 380 |
| 31. | Mapping Modifying and Tuning Tasks to Commands for HALDB Online Reorganization | 380 |
| 32. | Steps in Reorganizing a Database to Add a Logical Relationship. | 447 |
| 33. | Replace Rules for Logical View 1 | 479 |
| 34. | Replace Rules for Logical View 2 | 480 |
| 35. | Specifying Insert, Delete, and Replace Rules | 509 |
| 36. | Length and Format of an OSAM DEB. | 513 |

About This Book

This information is available as part of the Information Management Software for z/OS® Solutions Information Center. To view the information within the Information Management Software for z/OS Solutions Information Center, go to <http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp>. This information is also available in PDF and BookManager® formats. To get the most current versions of the PDF and BookManager formats, go to the IMS™ Library page at www.ibm.com/software/data/ims/library.html.

This book describes how to design, implement, and maintain different types of IMS databases and is divided into two parts:

- Part 1, “General Information on IMS Database Administration,” on page 1 describes important concepts to keep in mind throughout the database administration process.
- Part 2, “Administering IMS Databases,” on page 35 describes the steps in the database administration process.

With IMS Version 9, you can reorganize HALDB partitions online, either by using the integrated HALDB Online Reorganization function or by using an external product. In this information, the term HALDB Online Reorganization refers to the integrated HALDB Online Reorganization function that is part of IMS Version 9, unless otherwise indicated.

Prerequisite Knowledge

Before using this book, you should understand basic IMS concepts and your installation's IMS system. IMS can run in the following environments: DB Batch, DCCTL, TM Batch, DB/DC, DBCTL. You should understand the environments that apply to your installation. The IMS concepts explained here pertain only to administering the IMS database. You should know how to use DL/I calls and languages such as assembler, COBOL, PL/I, and C.

IMS Version 9: Application Programming: Design Guide describes how to design and code an application program.

For definitions of terms used in this manual and references to related information in other IMS manuals, see *IMS Version 9: Master Index and Glossary*.

IBM Product Names Used in This Information

In this information, the licensed programs shown in Table 1 are referred to by their short names.

Table 1. Licensed Program Full Names and Short Names

| Licensed program full name | Licensed program short name |
|---|-----------------------------|
| IBM® Application Recovery Tool for IMS and DB2® | Application Recovery Tool |
| IBM CICS® Transaction Server for OS/390® | CICS |
| IBM CICS Transaction Server for z/OS | CICS |
| IBM DB2 Universal Database™ | DB2 Universal Database |
| IBM DB2 Universal Database for z/OS | DB2 UDB for z/OS |

Table 1. Licensed Program Full Names and Short Names (continued)

| Licensed program full name | Licensed program short name |
|--|--|
| IBM Enterprise COBOL for z/OS and OS/390 | Enterprise COBOL |
| IBM Enterprise PL/I for z/OS and OS/390 | Enterprise PL/I |
| IBM High Level Assembler for MVS™ & VM & VSE | High Level Assembler |
| IBM IMS Advanced ACB Generator | IMS Advanced ACB Generator |
| IBM IMS Batch Backout Manager | IMS Batch Backout Manager |
| IBM IMS Batch Terminal Simulator | IMS Batch Terminal Simulator |
| IBM IMS Buffer Pool Analyzer | IMS Buffer Pool Analyzer |
| IBM IMS Command Control Facility for z/OS | IMS Command Control Facility |
| IBM IMS Connect for z/OS | IMS Connect |
| IBM IMS Connector for Java™ | IMS Connector for Java |
| IBM IMS Database Control Suite | IMS Database Control Suite |
| IBM IMS Database Recovery Facility for z/OS | IMS Database Recovery Facility |
| IBM IMS Database Repair Facility | IMS Database Repair Facility |
| IBM IMS DataPropagator™ for z/OS | IMS DataPropagator |
| IBM IMS DEDB Fast Recovery | IMS DEDB Fast Recovery |
| IBM IMS Extended Terminal Option Support | IMS ETO Support |
| IBM IMS Fast Path Basic Tools | IMS Fast Path Basic Tools |
| IBM IMS Fast Path Online Tools | IMS Fast Path Online Tools |
| IBM IMS Hardware Data Compression-Extended | IMS Hardware Data Compression-Extended |
| IBM IMS High Availability Large Database (HALDB) Conversion Aid for z/OS | IBM IMS HALDB Conversion Aid |
| IBM IMS High Performance Change Accumulation Utility for z/OS | IMS High Performance Change Accumulation Utility |
| IBM IMS High Performance Load for z/OS | IMS HP Load |
| IBM IMS High Performance Pointer Checker for OS/390 | IMS HP Pointer Checker |
| IBM IMS High Performance Prefix Resolution for z/OS | IMS HP Prefix Resolution |
| IBM z/OS Language Environment | Language Environment |
| IBM Tivoli® NetView® for z/OS | Tivoli NetView for z/OS |
| IBM WebSphere® Application Server for z/OS and OS/390 | WebSphere Application Server for z/OS |
| IBM WebSphere MQ for z/OS | WebSphere MQ |
| IBM WebSphere Studio Application Developer Integration Edition | WebSphere Studio |
| IBM z/OS | z/OS |
| IBM z/OS C/C++ | C/C++ |

How to Read Syntax Diagrams

The following rules apply to the syntax diagrams that are used in this information:


- Read the syntax diagrams from left to right, from top to bottom, following the path of the line. The following conventions are used:
 - The >>--- symbol indicates the beginning of a syntax diagram.
 - The ---> symbol indicates that the syntax diagram is continued on the next line.
 - The >--- symbol indicates that a syntax diagram is continued from the previous line.
 - The --->< symbol indicates the end of a syntax diagram.
- Required items appear on the horizontal line (the main path).




- Optional items appear below the main path.




If an optional item appears above the main path, that item has no effect on the execution of the syntax element and is used only for readability.




- If you can choose from two or more items, they appear vertically, in a stack. If you *must* choose one of the items, one item of the stack appears on the main path.




If choosing one of the items is optional, the entire stack appears below the main path.



If one of the items is the default, it appears above the main path, and the remaining choices are shown below.



- An arrow returning to the left, above the main line, indicates an item that can be repeated.




If the repeat arrow contains a comma, you must separate repeated items with a comma.




A repeat arrow above a stack indicates that you can repeat the items in the stack.

- Sometimes a diagram must be split into fragments. The syntax fragment is shown separately from the main syntax diagram, but the contents of the fragment should be read as if they are on the main path of the diagram.



fragment-name:



- In IMS, a b symbol indicates one blank position.
- Keywords, and their minimum abbreviations if applicable, appear in uppercase. They must be spelled exactly as shown. Variables appear in all lowercase italic letters (for example, *column-name*). They represent user-supplied names or values.
- Separate keywords and parameters by at least one space if no intervening punctuation is shown in the diagram.
- Enter punctuation marks, parentheses, arithmetic operators, and other symbols, exactly as shown in the diagram.
- Footnotes are shown by a number in parentheses, for example (1).

How to Send Your Comments

Your feedback is important in helping us provide the most accurate and highest quality information. If you have any comments about this or any other IMS information, you can take one of the following actions:

- Go to the IMS Library page at www.ibm.com/software/data/ims/library.html and click the Library Feedback link, where you can enter and submit comments.
- Send your comments by e-mail to imspubs@us.ibm.com. Be sure to include the title, the part number of the title, the version of IMS, and, if applicable, the specific location of the text on which you are commenting (for example, a page number in the PDF or a heading in the Information Center).

Summary of Changes

This section summarizes the significant improvements or enhancements for Administration Guide: Database Manager and refers you to relevant sections of this book for more information.

Changes to the Current Edition of This Book for IMS Version 9

This edition, which is available in softcopy format only, includes technical and editorial changes.

Changes to This Book for IMS Version 9

This book contains new information on the following subjects:

- The topic “Fast Path Virtual Storage Option” on page 139 contains new information about DEDB multi-area structures.
- The topic “Opening and Preopening DEDB Areas” on page 115 contains new information about enhancements to your ability to control the opening of DEDB areas during restart and after an IRLM reconnect.
- A new topic, “Reorganizing HALDBs” on page 364, discusses reorganizing HALDBs online with the integrated HALDB Online Reorganization function. This topic also includes additional information about reorganizing HALDBs offline.
- The existing topic Appendix E, “HALDB Partition Definition utility,” on page 517 has been updated to reflect the improvements that have been made to the Partition Definition utility for IMS Version 9.

Library Changes for IMS Version 9

Changes to the IMS Library for IMS Version 9 include the addition of one title, a change of one title, organizational changes, and a major terminology change. Changes are indicated by a vertical bar (|) to the left of the changed text.

The IMS Version 9 information is now available in the Information Management Software for z/OS Solutions Information Center, which is available at <http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp>. The Information Management Software for z/OS Solutions Information Center provides a graphical user interface for centralized access to the product information for IMS, IMS Tools, DB2 Universal Database (UDB) for z/OS, DB2 Tools, and DB2 Query Management Facility (QMF™).

New and Revised Titles

The following list details the major changes to the IMS Version 9 library:

- *IMS Version 9: IMS Connect Guide and Reference*

The library includes new information: *IMS Version 9: IMS Connect Guide and Reference*. This information is available in softcopy format only, as part of the Information Management Software for z/OS Solutions Information Center, and in PDF and BookManager formats.

IMS Version 9 provides an integrated IMS Connect function, which offers a functional replacement for the IMS Connect tool (program number 5655-K52). In this information, the term *IMS Connect* refers to the integrated IMS Connect function that is part of IMS Version 9, unless otherwise indicated.

- The information formerly titled *IMS Version 8: IMS Java User's Guide* is now titled *IMS Version 9: IMS Java Guide and Reference*. This information is available in softcopy format only, as part of the Information Management Software for z/OS Solutions Information Center, and in PDF and BookManager formats.
- To complement the IMS Version 9 library, a new book, *An Introduction to IMS* by Dean H. Meltz, Rick Long, Mark Harrington, Robert Hain, and Geoff Nicholls (ISBN # 0-13-185671-5), is available from IBM Press. Go to the IMS Web site at www.ibm.com/ims for details.

Organizational Changes

Organization changes to the IMS Version 9 library include changes to:

- *IMS Version 9: Customization Guide*
- *IMS Version 9: IMS Java Guide and Reference*
- *IMS Version 9: Messages and Codes, Volume 1*
- *IMS Version 9: Utilities Reference: System*

A new appendix has been added to the *IMS Version 9: Customization Guide* that describes the contents of the ADFSSMPL (also known as SDFSSMPL) data set.

The chapter titled "DLIModel Utility" has moved from *IMS Version 9: IMS Java Guide and Reference* to *IMS Version 9: Utilities Reference: System*.

The DLIModel utility messages that were in *IMS Version 9: IMS Java Guide and Reference* have moved to *IMS Version 9: Messages and Codes, Volume 1*.

Terminology Changes

IMS Version 9 introduces new terminology for IMS commands:

type-1 command

A command, generally preceded by a leading slash character, that can be entered from any valid IMS command source. In IMS Version 8, these commands were called *classic* commands.

type-2 command

A command that is entered only through the OM API. Type-2 commands are more flexible than type-1 commands and can have a broader scope. In IMS Version 8, these commands were called *IMSplex* commands or *enhanced* commands.

Accessibility Enhancements

Accessibility features help users who have physical disabilities, such as restricted mobility or limited vision, to use software products. The major accessibility features in z/OS products, including IMS, enable users to:

- Use assistive technologies such as screen readers and screen magnifier software
- Operate specific or equivalent features using only the keyboard
- Customize display attributes such as color, contrast, and font size

User Assistive Technologies

Assistive technology products, such as screen readers, work in conjunction with the IMS user interfaces. For information about how to use assistive technology to access these IMS user interfaces, consult the documentation of the assistive technology product.

Accessible Information

The IMS Version 9 product information is available online in the following accessible formats:

- HTML, in which you can access the IMS Version 9 library. You can find the IMS Version 9 library in the Information Management Software for z/OS Solutions Information Center at <http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp>.
- ISPF help panels, in which you can access the online help for the IMS Version 9 ISPF applications.

You can use screen readers and other assistive technologies to access the IMS Version 9 product information.

Keyboard Navigation of the User Interface

You can access the information center and IMS ISPF panel functions by using a keyboard or keyboard shortcut keys.

You can find information about navigating the information center using a keyboard in the information center home at <http://publib.boulder.ibm.com/infocenter/dzichelp/v2r2/index.jsp>.

For information about navigating the IMS ISPF panels using TSO/E or ISPF, refer to the *z/OS V1R1.0 TSO/E Primer*, the *z/OS V1R5.0 TSO/E User's Guide*, and the *z/OS V1R5.0 ISPF User's Guide, Volume 1*. These guides describe how to navigate each interface, including the use of keyboard shortcuts or function keys (PF keys). Each guide includes the default settings for the PF keys and explains how to modify their functions.

Part 1. General Information on IMS Database Administration

| | |
|--|----|
| Chapter 1. Introduction to IMS Databases | 3 |
| Database Administration Overview | 3 |
| DL/I | 3 |
| CICS | 3 |
| DBCTL and DCCTL | 4 |
| Open Database Access (ODBA) | 4 |
| Database Administration Tasks | 4 |
| Concepts and Terminology | 6 |
| How Data Is Stored in a Database | 6 |
| Root Segment | 7 |
| Parent and Child Segment | 7 |
| Segment Type and Occurrence | 7 |
| Relationship Between Segments | 8 |
| The Hierarchy | 8 |
| Numbering Sequence in a Hierarchy: Top to Bottom | 8 |
| Numbering Sequence in a Hierarchy: Movement and Position | 10 |
| Numbering Sequence in a Hierarchy: Level | 11 |
| The Database | 11 |
| The Database Record | 12 |
| The Segment | 14 |
| Segment Code | 14 |
| Delete Byte | 15 |
| Pointer and Counter Area | 15 |
| The Data Portion | 15 |
| The Three Data Portion Field Types | 15 |
| Optional Functions | 17 |
| How to Define Your Database to IMS | 18 |
| How Application Programs View the Database | 18 |
| Chapter 2. Standards and Procedures | 19 |
| Establishing Standards and Procedures | 19 |
| Naming Conventions | 21 |
| General Rules for Establishing Naming Conventions | 21 |
| HALDB Naming Conventions | 22 |
| HALDB Partition Names | 23 |
| HALDB DD names | 23 |
| HALDB Data Set Names | 23 |
| Chapter 3. Review Process | 25 |
| The Design Review | 25 |
| Role of the Database Administrator | 25 |
| General Information about Reviews | 25 |
| Design Review 1 | 26 |
| Design Review 2 | 26 |
| Design Review 3 | 27 |
| Design Review 4 | 27 |
| Code Inspection 1 | 28 |
| Who Attends Code Inspection 1 | 28 |
| Code Inspection 2 | 28 |
| Security Inspection | 29 |
| Post-Implementation Review | 29 |
| Chapter 4. Security | 31 |

| | |
|---|----|
| Restricting the Scope of Data Access | 31 |
| Restricting Processing Authority. | 31 |
| Restricting Access by Non-IMS Programs | 33 |
| Protecting Data with VSAM Passwords | 33 |
| Encrypting Your Database | 33 |
| Using the Dictionary to Help Establish Security | 33 |

Chapter 1. Introduction to IMS Databases

This chapter describes the tasks of database administration and discusses the key concepts and terms used when administering IMS Database Manager.

In this Chapter:

- “Database Administration Overview”
- “Open Database Access (ODBA)” on page 4
- “Database Administration Tasks” on page 4
- “Concepts and Terminology” on page 6
- “Optional Functions” on page 17
- “How to Define Your Database to IMS” on page 18
- “How Application Programs View the Database” on page 18

Database Administration Overview

The task of database administration is to design, implement, and maintain databases. This book describes the tasks involved in administering the Information Management System Database Manager (IMS DB). IMS is composed of two parts: IMS Database Manager and IMS Transaction Manager. IMS Database Manager manages the physical storage of records in the database. IMS Transaction Manager manages the terminal network, the input and output of messages, and online system resources. The administration of IMS Transaction Manager is covered in the *IMS Version 9: Administration Guide: System* and the *IMS Version 9: Administration Guide: Transaction Manager*.

This book presents the database administration tasks in the order in which you normally perform the tasks. You perform some tasks in a specific sequence in the database development process while other tasks are ongoing. It is important for you to grasp not only what the tasks are (see “Database Administration Tasks” on page 4), but also how they interrelate.

This first part of the book provides important concepts and procedures for the entire database administration process. The second part contains the chapters corresponding to particular tasks of database administration.

DL/I

Data Language/I (DL/I) is the IMS data manipulation language, which is a common high-level interface between a user application and IMS. DL/I calls are invoked from application programs written in languages such as PL/I, COBOL, VS Pascal, C, and Ada. It can also be invoked from assembler language application programs by subroutine calls. IMS lets the user define data structures, relate structures to the application, load structures, and reorganize structures.

Related Reading: For detailed information about how application programs use DL/I, see *IMS Version 9: Application Programming: Database Manager* and *IMS Version 9: Application Programming: EXEC DLI Commands for CICS and IMS*.

CICS

Customer Information Control System (CICS) accesses IMS databases through the database resource adapter (DRA). CICS or other transaction management

Database Administration Overview

subsystems (excluding IMS Transaction Manager) can access IMS full-function databases and data entry databases (DEDBs) in a DB/DC or DBCTL environment through the DRA.

Whenever tasks differ for CICS users, a brief description about the differences is included.

DBCTL and DCCTL

Database Control (DBCTL) supports non-message-driven batch message processing (BMP) programs. DBCTL has its own log and participates in database recovery. Locking is provided by IMS program isolation (PI) or the internal resource lock manager (IRLM).

Data Communications Control (DCCTL) is a transaction management subsystem that does not support full-function DEDBs or MSDBs (main storage databases), but does support GSAM databases in BMP regions. To access databases in a DCCTL environment, DCCTL must connect to an external subsystem that provides database support.

Open Database Access (ODBA)

Any program that runs in a z/OS address space can access IMS DB through the Open Database Access (ODBA) callable interface. Any z/OS application program running in a z/OS address space that is managed by the z/OS Resource Recovery Service (RRS/MVS) can access IMS full-function databases and data entry databases (DEDBs). z/OS application programs that use the ODBA interface are called ODBA applications.

Related Reading: For a description of RRS and its uses, see the information on RRS Distributed Sync Point in *IMS Version 9: Administration Guide: Transaction Manager*.

From the perspective of IMS, the z/OS address space involved appears to be another region called the z/OS application region.

Types of programs that can call the ODBA interface include:

- DB2 for z/OS stored procedures, including COBOL, PL/I, and Java procedures
- WebSphere for z/OS and OS/390 Enterprise Java Beans
- Other z/OS applications

Database Administration Tasks

Participating in design reviews. Design reviews are a series of formal meetings you attend in which the design and implementation of the database are examined. Design reviews are an ongoing task during the design and implementation of a database system. They are also held when new applications are added to an existing system.

Analyzing data requirements. After the users at your installation identify their data processing requirements, you will construct data structures. These structures show what data will be in your database and how it will be organized. This task precedes the actual design of the database.

Designing your database. After data structures are identified, the next step is to design your database. Database design involves:

- Choosing how to physically organize your data

- Deciding which IMS processing options you need to use
- Making a series of decisions about design that determine how well your database performs and uses available space

Developing a test database. Before the applications that will use your database are cut over to production status, they should be tested. Depending on the form of your existing data, you can use one or more of the IMS Database Design Aids to design, create, load, and test your test database.

Implementing your database design. After your database is designed, implement the design by describing the database's characteristics and how application programs will use it to IMS. This task consists of coding database descriptions (DBDs) and program specification blocks (PSBs), both of which are a series of macro statements. Another part of implementing the database design is determining whether to have the application control blocks (ACBs) of the database prebuilt or built dynamically.

Loading your database. After database characteristics are defined, write an initial load program to put your data into the database. After you load the database, application programs can be run against it.

Monitoring your database. When the database is running, routinely monitor its performance. A variety of tools for monitoring the IMS system are available.

Tuning your database. Tune your database when performance degrades or utilization of external storage is not optimum. Routine monitoring helps you determine when the system needs to be tuned and what type of tuning needs to be done. Like monitoring, the task of tuning the database is ongoing.

Modifying your database. As new applications are developed or the needs of your users change, you might need to make changes to your database. For example, you can change database organization, database hierarchies (or the segments and fields within them), and you can add or delete one or more partitions. Like monitoring and tuning, the task of modifying the database is ongoing.

Recovering your database. Database recovery involves restoring a database to its original condition after it is rendered invalid by some failure. The task of developing recovery procedures and performing recovery is an important one. However, because it is difficult to separate data recovery from system recovery, the task of recovery is treated separately in *IMS Version 9: Operations Guide*.

You can use Database Recovery Control (DBRC) in recovering your databases. If your databases are registered in RECON, DBRC gains control during execution of these IMS utilities:

- Database Image Copy
- Online Database Image Copy
- Database Image Copy 2
- Change Accumulation
- Database Recovery
- Log Recovery
- Log Archive
- DEDB area data set create
- HD and HISAM Reorganization Unload and Reload
- HALDB Index/ILDS Rebuild

You must ensure that all database recoveries use the current IMS utilities, rather than those of earlier releases.

Related Reading: For more information on using these database utilities, see the *IMS Version 9: Utilities Reference: System* and the *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

Establishing security. You can keep unauthorized persons from accessing the data in your database by using program communication blocks (PCBs). With PCBs, you can control how much of the database a given user can see, and what can be done with that data. In addition, you can take steps to keep non-IMS programs from accessing your database.

Setting up standards and procedures. It is important to set standards and procedures for application and database development. This is especially true in an environment with multiple applications. If you have guidelines and standards, you will save time in application development and avoid problems later on such as inconsistent naming conventions or programming standards.

Concepts and Terminology

This topic discusses the terms and concepts you need to understand to perform the administration tasks just outlined.

To understand this topic, you must know what a DL/I call is and how to code it. You must understand function codes and Segment Search Arguments (SSAs) in DL/I calls and know what is meant when a call is referred to as qualified or unqualified (explained in *IMS Version 9: Application Programming: Database Manager*).

How Data Is Stored in a Database

The data in a database is grouped into a series of *database records*. Each database record is composed of smaller groups of data called *segments*. A segment is the smallest piece of data IMS can store. Segments, in turn, are made up of one or more *fields*.

Figure 1 on page 7 shows a record in a school database. Each of the boxes is a segment or separate group of data in the database record. The segments in the database record contain the following information:

| | |
|----------------|---|
| COURSE | The name of the course |
| INSTR | The name of the teacher of the course |
| REPORT | A report the teacher needs at the end of the course |
| STUDENT | The names of students in the course |
| GRADE | The grade a student received in the course |
| PLACE | The room in which the course is taught |

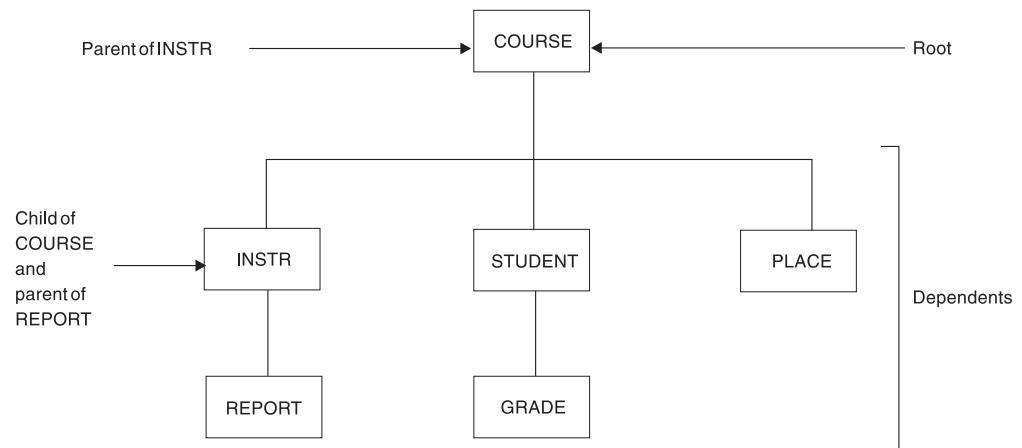


Figure 1. Segment Types in the School Database Record

The segments within a database record exist in a *hierarchy*. A hierarchy is the order in which segments are arranged. The order implies something. The school database is storing data about courses that are taught. The COURSE segment is at the top of the hierarchy. The other types of data in segments in the database record would be meaningless if there was no COURSE.

Root Segment

The COURSE segment is called the *root segment*. Only one root segment exists within a database record. All other segments in the database record (such as: INSTR, REPORT, STUDENT, GRADE, and PLACE) are called *dependent segments*. The existence of dependent segments hinges on the existence of a root segment. For example, without the root segment COURSE, there would be no reason for having a PLACE segment stating in which room the course was held.

The third level of dependent segments, REPORT and GRADE, is subject to the existence of second level segments INSTR and STUDENT. For example, without the second level segment STUDENT, there would be no reason for having a GRADE segment indicating the grade the student received in the course.

Parent and Child Segment

Another set of words used to refer to how segments relate to each other in a hierarchy is *parent segment* and *child segment*. A parent segment is any segment that has a dependent segment beneath it in the hierarchy. COURSE is the parent of INSTR, and INSTR is the parent of REPORT. A child segment is any segment that is a dependent of another segment above it in the hierarchy. REPORT is the child of INSTR, and INSTR is the child of COURSE. Note that INSTR is both a parent segment in its relationship to REPORT and a child segment in its relationship to COURSE.

Segment Type and Occurrence

The terms used to describe segments thus far (root, dependent, parent, and child) describe the *relationship* between segments. The terms *segment type* and *segment occurrence*, however, distinguish between a type of segment in the database (the COURSE segment or the INSTR segment) and a specific segment (the course segment for a math course).

The previous database is actually the design of the database. It shows the segment types for the database. Figure 2 on page 8 shows the actual database record with the segment occurrences.

Concepts and Terminology

A segment occurrence is a single specific segment. Math is a single occurrence of the COURSE segment type. Baker and Coe are multiple occurrences of the STUDENT segment type.

Relationship Between Segments

One final term for describing segments is *twin segment*. Twin (like root, dependent, parent, and child) describes a relationship between segments. Twin segments are multiple occurrences of the same segment type under a single parent. In Figure 2, the segments Baker and Coe are twins. They have the same parent (Math), and are of the same segment type (STUDENT). Pass and Inc are not twins. Although Pass and Inc are the same segment type (GRADE), they do not have the same parent. Pass is the child segment of Baker, and Inc is the child segment of Coe.

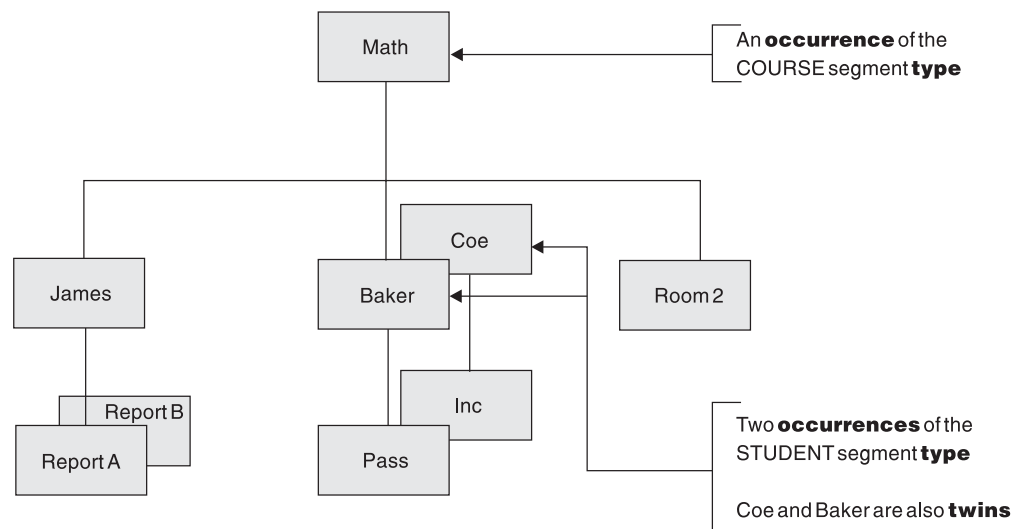


Figure 2. Segment Occurrences in a School Database Record

The following topic discusses the hierarchy in more detail. Subsequent topics describe the objects in a database, what they consist of and the rules governing their existence and use. These objects are:

- The database record
- The segments in a database record
- The fields within a segment

The Hierarchy

A database is composed of a series of database records, records contain segments, and segments are arranged in a hierarchy in the database record.

Numbering Sequence in a Hierarchy: Top to Bottom

When a database record is stored in the database, the hierarchic arrangement of segments in the database record is the order in which segments are stored. Starting at the top of a database record (at the root segment), segments are stored in the database in the sequence shown by the numbers in Figure 3 on page 9.

The sequence goes from the top of the hierarchy to the bottom in the first (left most) *path* or leg of the hierarchy. When the bottom of the database is reached, the sequence is from left to right. When all segments have been stored in that path of the hierarchy, the sequencing begins in the next path to the right, again proceeding

from top to bottom and then left to right. (In the second leg of the hierarchy there is nothing to go to at the right.) The sequence in which segments are stored is loosely called “*top to bottom, left to right*.”

Figure 3 shows sequencing of segment types for the school database shown in Figure 1 on page 7. The sequence of segment types are stored in the following order:

1. COURSE (top to bottom)
2. INSTR
3. REPORT
4. STUDENT (left to right)
5. GRADE (top to bottom)
6. PLACE (left to right)

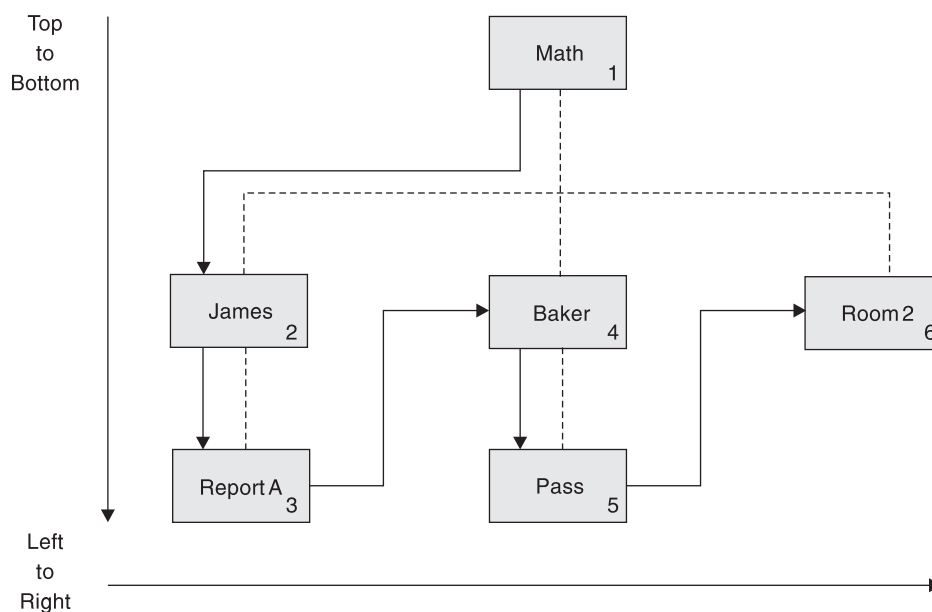


Figure 3. Hierarchic Sequence of Segment Types for School Database

Figure 4 on page 10 shows the segment occurrences for the school database record as shown in Figure 2 on page 8. Because there are multiple occurrences of segment types, segments are read “front to back” in addition to “top to bottom, left to right.” The segment occurrences for the school database are stored in the following order:

1. Math (top to bottom)
2. James
3. ReportA
4. ReportB (front to back)
5. Baker (left to right)
6. Pass (top to bottom)
7. Coe (front to back)
8. Inc (top to bottom)
9. Room2 (left to right)

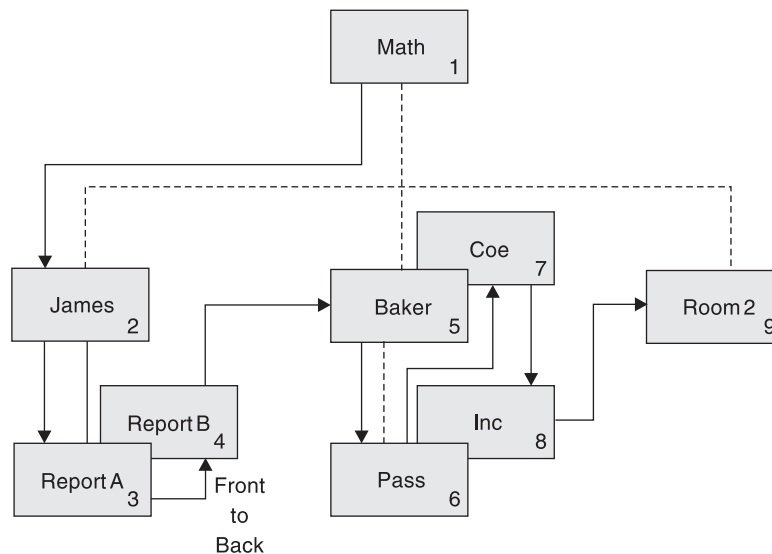


Figure 4. Hierarchic Sequence of Segment Occurrences for School Database

Note that the numbering sequence is still initially from top to bottom. At the bottom of the hierarchy, however, observe that there are two occurrences of the REPORT segment.

Because you are at the bottom of the hierarchy, both segment occurrences are picked up before you move to the right in this path of the hierarchy. Both reports relate to the instructor segment James; therefore it makes sense to keep them stored together in the database. In the second path of the hierarchy, there are also two segment occurrences in the student segment. You are not at the bottom of the hierarchic path until you reach the grade segment Pass. Therefore, sequencing is not “interrupted” by the two occurrences of the student segment Baker and Coe. This makes sense because you are keeping student and grade Baker and Pass together.

Note that the grade Inc under student Coe is not considered another occurrence under Baker. Coe and Inc become a separate path in the hierarchy. Only when you reach the bottom of a hierarchic path is the “top to bottom, left to right” sequencing interrupted to pick up multiple segment occurrences. You can refer to sequencing in the hierarchy as “top to bottom, front to back, left to right”, but “front to back” only occurs at the bottom of the hierarchy. Multiple occurrences of a segment at any other level are sequenced as separate paths in the hierarchy.

As noted before, this numbering of segments represents the sequence in which segments are stored in the database. If an application program requests all segments in a database record in hierarchic sequence or issues Get-Next (GN) calls, this is the order in which segments would be presented to the application program.

Numbering Sequence in a Hierarchy: Movement and Position

Other terms that show the numbering sequence in a hierarchy are: *movement* and *position*. When talking about movement through the hierarchy, it always means moving in the sequence implied by the numbering scheme. Movement can be forward or backward. When talking about position in the hierarchy, it means being located (positioned) at a specific segment. The terms movement and position are used when talking about how segments are accessed when an application program issues a call.

A segment is the smallest piece of data IMS can store. If an application program issues a Get-Unique (GU) call for the student segment BAKER (see Figure 4 on page 10), the current position is immediately after the BAKER segment occurrence. If an application program then issues an unqualified GN call, IMS moves forward in the database and returns the PASS segment occurrence.

Numbering Sequence in a Hierarchy: Level

A final term you need to know about hierarchies is: *level*. Level is the position of a segment in the hierarchy in relation to the root segment. The root segment is always on level one. Figure 5 illustrates levels of the database record shown in Figure 2 on page 8.

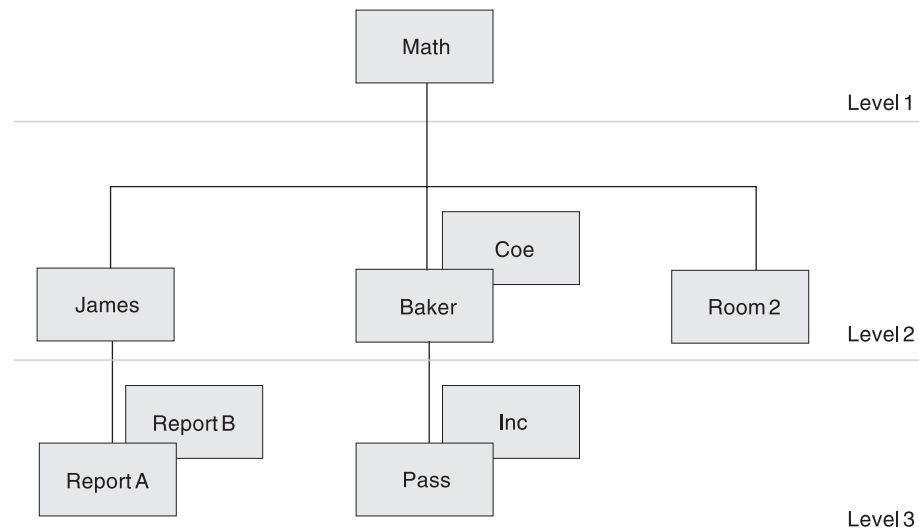


Figure 5. Levels in the Database

The Database

IMS allows you to define many different database types. You define the database type that best suits your application's processing requirements. You need to know that each IMS database has its own access method, because IMS runs under control of the z/OS operating system. The operating system does not know what a segment is because it processes logical records, not segments. IMS access methods therefore manipulate segments in a database record. When a logical record needs to be read, operating system access methods (or IMS) are used.

Table 2 lists the IMS database types you can define, the IMS access methods they use and the operating system access methods you can use with them. Although each type of database varies slightly in its access method, they all use database records.

Table 2. Types of IMS Databases and the z/OS Access Methods They Can Use

| Type of IMS Database | Full Name of Database Type | IMS or Operating System Access Methods that Can Be Used |
|----------------------|---|---|
| HSAM | Hierarchical Sequential Access Method | BSAM or QSAM |
| SHSAM | Simple Hierarchical Sequential Access Method | BSAM or QSAM |
| HISAM | Hierarchical Indexed Sequential Access Method | VSAM |

Concepts and Terminology

Table 2. Types of IMS Databases and the z/OS Access Methods They Can Use (continued)

| Type of IMS Database | Full Name of Database Type | IMS or Operating System Access Methods that Can Be Used |
|----------------------|---|---|
| SHISAM | Simple Hierarchical Indexed Sequential Access Method | VSAM |
| GSAM | Generalized Sequential Access Method | QSAM/BSAM or VSAM |
| HDAM | Hierarchical Direct Access Method | VSAM or OSAM |
| PHDAM | Partitioned Hierarchical Direct Access Method | VSAM or OSAM |
| HIDAM | Hierarchical Indexed Direct Access Method | VSAM or OSAM |
| PHIDAM | Partitioned Hierarchical Indexed Direct Access Method | VSAM or OSAM |
| DEDB ¹ | Data Entry Database | Media Manager |
| MSDB ² | Main Storage Database | N/A |

Notes:

1. For DBCTL, only available to BMPs
2. Not applicable to DBCTL

The Database Record

A database consists of a series of database records, and a database record consists of a series of segments. Another thing to understand is that a specific database can only contain one kind of database record. In the school database, for example, you can place as many school records as desired. You could not, however, create a different type of database record, such as the medical database record shown in Figure 6, and put it in the school database.


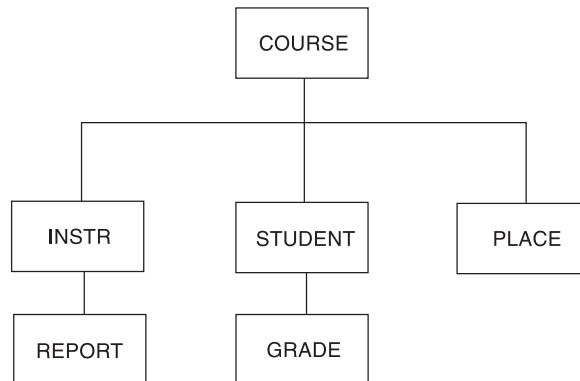



Figure 6. An Example of a Medical Database Record

The only other thing to understand is that a specific database record, when stored in the database, does not need to contain all the segment types you originally designed. To exist in a database, a database record need only contain an occurrence of the root segment. In the school database, all four of the records shown in Figure 7 on page 13 can be stored.


Database Record 1



Database Record 2



Database Record 3



Database Record 4




Figure 7. Example of Records That Can Be Stored in the School Database

However, no segment can be stored unless its parent is also stored. For example, you could *not* store the records shown in Figure 8.




Figure 8. Records that Cannot be Stored in the School Database

Occurrences of any of the segment types can later be added to or deleted from the database.

The Segment

A database record consists of one or more segments, and the segment is the smallest piece of data IMS can store. Here are some additional facts you need to know about segments:

- A database record can contain a maximum of 255 segment types. The space you allocate for the database limits the number of segment occurrences.
- You determine the length of a segment; however, a segment cannot be larger than the physical record length of the device on which it is stored.
- The length of segments is specified by segment type. A segment type can be either variable or fixed in length.

Figure 9 shows the format of a fixed-length segment. Figure 10 shows the format of a variable-length segment. Segments consist of two parts (a prefix and the data), except when using a SHSAM or SHISAM database. In SHSAM and SHISAM databases, the segment consists of only the data. In a GSAM database, segments do not exist (see “GSAM Databases” on page 80 for more information about GSAM databases).




Figure 9. Format of Fixed-Length Segments




Figure 10. Format of Variable-Length Segments

IMS uses the prefix portion of the segment to “manage” the segment. The prefix portion of a segment consists of: segment code, delete byte, and in some databases, a pointer and counter area. Application programs do not “see” the prefix portion of a segment. The data portion of a segment contains your data, arranged in one or more fields.

Related Reading: For information on MSDB and DEDB segments, see “Main Storage Databases (MSDBs)” on page 132 and “Data Entry Databases” on page 113.

Segment Code

IMS needs a way to identify each segment type stored in a database. It uses the segment code field for this purpose. When loading a segment type, IMS assigns it a unique identifier (an integer from 1 to 255). IMS assigns numbers in ascending sequence, starting with the root segment type (number 1) and continuing through all dependent segment types in hierarchic sequence.

Delete Byte

When an application program deletes a segment from a database, the space it occupies might or might not be immediately available to reuse. Deletion of a segment is described in the discussions of the individual database types. For now, know that IMS uses this prefix byte to track the status of a deleted segment.

Related Reading: For information on the meaning of each bit in the delete byte, see Appendix A, “Meaning of Bits in the Delete Byte,” on page 469.

Pointer and Counter Area

The pointer and counter area exists in HDAM, PHDAM, HIDAM, and PHIDAM databases, and, in some special circumstances, HISAM databases. The pointer and counter area can contain two types of information:

- Pointer information consists of one or more addresses of segments to which a segment points.
- Counter information is used when logical relationships, an optional function of IMS, are defined.

The length of the pointer and counter area depends on how many addresses a segment contains and whether logical relationships are used. These topics are covered in more detail later in this book.

The Data Portion

The data portion of a segment contains one or more data elements. The data is processed and unlike the prefix portion of the segment, seen by an application program.

The application program accesses segments in a database using the name of the segment type. If an application program needs to reference part of a segment, a field name can be defined to IMS for that part of the segment. Field names are used in segment search arguments (SSAs) to qualify calls. An application program can see data even if you do not define it as a field. But an application program cannot qualify an SSA on the data unless it is defined as a field.

The maximum number of fields that you can define for a segment type is 255. The maximum number of fields that can be defined for a database is 1000. Note that 1000 refers to *types* of fields in a database, not occurrences. The number of occurrences of fields in a database is limited only by the amount of storage you have defined for your database.

The Three Data Portion Field Types

You can define three field types in the data portion of a segment: a sequence field, data fields, and for variable-length segments, a size field stating the length of the segment. The first two field types contain your data, and an application program can use both to qualify its calls. However, the sequence field has some other uses besides that of containing your data.

You can use a sequence field, often referred to as a key, to keep occurrences of a segment type in key sequence under a given parent. For example, in the database record shown in Figure 11 on page 16, there are three segment occurrences of the STUDENT segment, and the STUDENT segment has three data elements.

Concepts and Terminology

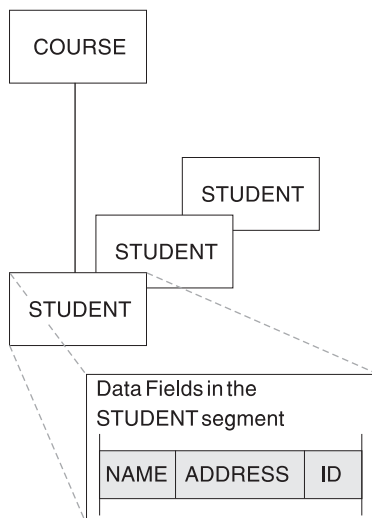


Figure 11. Three Segment Occurrences and Three Data Elements of the STUDENT Segment

Suppose you need the STUDENT segment, when stored in the database, to be in alphabetic order by student name. If you define a field on the NAME data as a *unique* sequence field, IMS stores STUDENT segment occurrences in alphabetical sequence. Figure 12 shows three occurrences of the STUDENT segment in alphabetical sequence.




Figure 12. Example of STUDENT Segments Stored in Alphabetical Order

When you define a sequence field in a root segment of a HISAM, HDAM, PHDAM, HIDAM, or PHIDAM database, an application program can use it to access a specific root segment, and thus a specific database record. By using a sequence field, an application program does not need to search the database sequentially to find a specific database record, but can retrieve records sequentially (for HISAM, HIDAM, and PHIDAM databases).

You can also use a sequence field in other ways when using the IMS optional functions of logical relationships or secondary indexing. These other uses are discussed in detail later in this book.

The important things to know now about sequence fields are that:

- You do not always need to define a sequence field. This book describes cases where a sequence field is necessary.
- The sequence field value can be defined as unique or non-unique.
- The data or value in the sequence field is called the “key” of the segment.

Optional Functions

IMS has several optional functions you can use for your database. These are discussed briefly below and described in detail in Chapter 8, “Choosing Optional Database Functions,” on page 157. You need a cursory understanding of these functions before reading this book because they may be referred to before they are actually described.

The functions are as follows:

Logical relationships is a function you can use to let an application program access a logical database record. A logical database record can consist of segments from one or more physical database records. Physical database records can be stored in one or more databases. Thus, a logical database record lets an application program view a database structure that is different from the physical database structure.

For example, if a logical data structure contains segments from two different physical databases, a segment can be accessed from two different paths:

- A segment can be physically stored in the path where it is most frequently used and where the most urgent response time is required.
- A pointer containing the location of the segment can be physically stored in the alternate path needed by another application program.

Secondary indexing is a function you can use to access segments in a database in a sequence other than the one defined in the sequence field.

Variable-length segments is a function you can use to make the data portion of a segment type variable in length. Use variable-length segments when the size of the data portion of a segment type varies greatly from one segment occurrence to the next. With variable-length segments, you define the minimum and maximum length of a segment type. Defining both minimum and maximum length saves space in the database whenever a segment is shorter than the maximum length.

Field-level sensitivity is a function you can use to:

- Deny an application program access to selected fields in a segment for security purposes.
- Allow an application program to use a subset of the fields that make up a segment (and not process fields it does not use) or use fields in a segment in a different order. Use field-level sensitivity in this way to accommodate the differing needs of your application programs.

Segment edit/compression is a function you can use with segments to:

- Encode or “scramble” segment data when it is on the device so only application programs with access to the segment receive the data in decoded form.
- Edit data so application programs can receive data in a format other than the one in which it is stored.
- Compress data when writing a segment to the device, so the Direct Access Storage Device (DASD) is better used.

A *Data Capture exit routine* is used to capture segment data when an application program updates IMS databases with an insert, replace, or delete call. This is a synchronous activity that happens within the unit of work or application update. Captured data is used for data propagation to DB2 UDB for z/OS databases. You can also use Data Capture exit routines to perform tasks other than data propagation.

Optional Functions

Asynchronous Data Capture is a function you use to capture segment data when an application program updates IMS databases with an insert, replace, or delete call. This is an asynchronous activity that happens outside of the unit of work or application update. Captured data is used for data propagation to DB2 UDB for z/OS databases asynchronously. You can also use Asynchronous Data Capture to perform tasks other than data propagation.

IMS DataPropagator allows you to propagate the changed data to or from IMS and DB2 UDB for z/OS both synchronously and asynchronously.

Related Reading: For more information on IMS DataPropagator see *IMS DataPropagator for z/OS: An Introduction*.

Multiple data set groups is a function you can use to put some segments in a database record in data sets other than the primary data set. This can be done without destroying the hierarchic sequence of segments in a database record.

One reason to use multiple data set groups is to accommodate the differing needs of your applications. By using multiple data set groups, you can give an application program fast access to the segments in which it is interested. The application program simply bypasses the data sets containing unnecessary segments. Another reason for using multiple data set groups is to improve performance by, for example, separating high-use segments from low-use segments. You might also use multiple data set groups to save space by putting segment types whose size varies greatly from the average in a separate data set group.

How to Define Your Database to IMS

Define the characteristics of your database to IMS by coding and generating a DBD (database description). A DBD is a series of macro instructions that describes a database's organization and access method, the segments and fields in a database record, and the relationship between types of segments.

If you have the IBM DB/DC (database/data communication) Data Dictionary, you can use it to define your database (except for DEDBs and MSDBs). The DB/DC Data Dictionary may contain all the information you need to produce a DBD.

How Application Programs View the Database

You control how an application program views your database. An application program might not need use of all the segments or fields in a database record. And an application program may not need access to specific segments for security or integrity purposes. An application program may not need to perform certain types of operations on some segments or fields. For example, an application program needs read access to a SALARY segment but not update access. You control which segments and fields an application can view and which operations it can perform on a segment by coding and generating a PSB (program specification block).

A PSB is a series of macro instructions that describe an application program's access to segments in the database. A PSB consists of one or more program communication blocks (PCB), and each PCB describes the application program's ability to read and use the database. For example, an application program can have different views and uses of the same database. An application program can access several different databases and can have several PCBs in its PSB.

If you have the IBM DB/DC Data Dictionary, you can use it to define an application program's access to the database. It can contain all the information needed to produce a PSB.

Chapter 2. Standards and Procedures

This chapter examines the following areas:

- “Establishing Standards and Procedures”
- “Naming Conventions” on page 21

Establishing Standards and Procedures

You should give careful thought to developing standards and procedures for your database system. Providing standards and procedures results in:

- Improved quality of application systems (because setting up and following standards and procedures gives you greater control over your entire application development process)
- Improved productivity in application and database design (because guidelines for design decisions exist)
- Improved productivity of application coding (because coding standards and procedures exist)
- Better communication between you and application developers (because both of you have clearly defined responsibilities)
- Improved reliability and recoverability in operations (because you have clear and well-understood operating procedures)

You must set up and test procedures and standards for database design, application development, application programs’ use of the database, application design, and for batch operation. These standards are guidelines that change when installation requirements change.

In the area of database design, for example, you can establish standard practices for handling the following items:

- Database structure and segmentation
 - Number of segments within a database
 - Placement of segments
 - Size of segments
 - Use of variable-length segments
 - When to use segment edit/compression
 - When to use secondary data set groups
 - Number of databases within an application
 - When and how to use field-level sensitivity
 - Database size
- Access methods
 - When to use HISAM
 - Choice of record size for HISAM
 - HISAM organization using VSAM
 - When to use GSAM
 - Use of physical child/physical twin pointers
 - Use of twin backward pointers
 - Use of child last pointers
 - HIDAM or PHIDAM index organization using VSAM

Establishing Standards and Procedures

- HIDAM or PHIDAM pointer options at the root level
- Sequencing twin chains
- Use of HD free space
- When to use HDAM or PHDAM
- Processing an HDAM or a PHDAM database sequentially
- Use of the “byte limit count” for HDAM or PHDAM
- Use of twin backward pointer for HDAM or PHDAM roots
- Use of free space with HDAM or PHDAM
- When to use DEDBs
- Processing DEDBs sequentially
- Use of DEDB parameters
- Use of subset pointers
- Use of multiple area data sets
- Secondary indexing
 - For sequential processing
 - On volatile segments
 - In HISAM databases
 - Use of unique secondary indexes
 - Use of sparse indexing
 - Processing of the secondary index as a separate database
- Logical relationships
 - Use of direct pointers versus symbolic pointers
 - Avoidance of long logical twin chains
 - Sequencing of the logical twin chain
 - Placement of the real logical child segment

In the area of application programs use of the database, establish standards for the following:

- Putting update and read functions in separate programs
- How many transaction types to allow per application program
- When applications are to issue a deliberate abnormal termination and the range of abend codes permitted applications
- Whether application programs are permitted to issue messages to the master terminal
- The method of referencing data in the IOAREA, and referencing IMS variables (such as PCBs and SSAs)
- Use of predefined structures (PCB masks, SSAs, or database segment formats) by applications
- Use of GU calls to the message queue
- Re-usability of MPP and BMP programs
- Use of qualified calls and SSAs
- Use of path calls
- Use of the CHANGE call
- Use of the “system” calls (PURG, LOG, STAT, SNAP, GCMD, and CMD)

In the area of application design, establish procedures to govern the following:

- The interaction between you and the application designer

- Use of the dictionary or COPY or STRUCTURE libraries for data elements and structures
- The holding of design reviews and inspections

In the area of batch operations, you can consider developing:

- Procedures to limit access to computer facilities
- A control point to ensure that:
 - Jobs contain complete and proper submittal documentation
 - Jobs are executed successfully on schedule
 - Correct input/output volumes are used, and output is properly distributed
 - Test programs are executed only in accordance with a defined test plan
 - An incident report is maintained to ensure all problems are recorded and reported to the responsible parties
- Normal operating procedures. These operating procedures include operations schedules, cold start, warm start, shutdown procedures, and scheduling and execution of batch programs.
- Procedures for emergency situations. During an emergency, the environment is one of stress. Documented procedures provide step-by-step guidance to resolve such situations. These procedures should include emergency restart, database backout, database recovery, log recovery, and batch program restart.

Related Reading: For a more complete treatment of recovery procedures, see *IMS Version 9: Operations Guide*.

- A master terminal operator's guide for the installation. This guide should be supplemented by *IMS Version 9: Command Reference*.
- A master operations log. This log could contain a record of system availability, time and type of failure, and cause of the failure, recovery steps taken, and type of system termination if normal.
- A system maintenance log. This log could contain a record of all release and modification levels, release dependencies, program temporary fixes (PTFs) applied, status of APARs and date submitted, and bypass solutions.

Naming Conventions

This topic contains information about:

- "General Rules for Establishing Naming Conventions"
- "HALDB Naming Conventions" on page 22

General Rules for Establishing Naming Conventions

Good naming conventions are mandatory in a data processing project, especially in an environment with multiple applications. Some general rules to follow in setting up naming conventions are:

- Each name should be unique.
- Each name should be meaningful and identifiable. You should be able to identify the type of thing being referred to by its name.

Table 3 on page 22 is an example of minimal naming conventions. They are presented only as an example, and you can establish your own naming conventions.

Establishing Naming Conventions

Table 3. Example of Naming Conventions

| CATEGORY | CONVENTION |
|-------------|---|
| SYSTEM | S as first letter |
| JOB | J as first letter |
| PROGRAM | P as first letter if IMS program (to match PSB) G as first letter otherwise |
| MODULE | M as first letter |
| COPY | C as first letter for member containing the segment structure A as first letter for member containing all the SSAs for the segment Remainder must be the same as the segment name |
| TRANSACTION | T as first letter |
| PSB | P as first letter |
| PCB | Same name as PSB Note: Occurrence number indicates position in PSB |
| DATABASE | Dtaaann Where Indicates t Database type. The database can be one of the following types: P Physical L Logical X Primary index Y Secondary index aaa A unique database identifier common to all logical and index databases based on the same physical database nn A unique identifier, if there are multiple logical or secondary index databases |
| SEGMENT | Saaabbbb aaa A unique database identifier; same as the physical database in which the segment occurs Note: Concatenated segments should have an aaa value corresponding to the aaa of the logical child segment. bbbb An identifier for the user name R First letter for 'segments' that are non-DL/I file record definitions O First letter for any other data areas, for example, terminal I/O areas, control blocks, report lines, and so on. |
| ELEMENT | E as first letter |

HALDB Naming Conventions

Unique HALDB naming conventions are described in the following topics:

- Partition names
- DDNAMEs
- Data set names

HALDB Partition Names

Each HALDB partition name is 1 to 7 bytes in length and must be unique among the database names, DEDB names, Area names, and partition names in one RECON data set. The HALDB partition names are used to represent specific partitions and are used interchangeably with database names in commands.

HALDB DD names

IMS constructs the DD names for each partition by adding a 1-byte suffix to the partition name for the data sets in that partition. The suffix for the first DD name is A, the suffix for the second DD name is B, and so on up to J.

For a PSINDEX database, there is only one data set per partition, so only one DD name with a suffix of A is required.

The resulting DD names with the suffix might match already existing DD names and you must avoid duplication of DD names. The DD names are not case sensitive and can result in JCL errors if specified in lower case in batch jobs.

Table 4 shows the suffixes for the different DD names.

Table 4. Suffixes for DD names

| Corresponding Data Set | Suffix | HALDB OLR Suffix |
|-------------------------------|--|--|
| Primary data area | A (first data set) through J (last data set) | M (first data set) through V (last data set) |
| Primary index (PHIDAM only) | X | Y |
| Indirect list data set (ILDS) | L | L (OLR uses the same ILDS) |

Extended Naming Convention for DD Names When Using HALDB Online

Reorganization: When you reorganize either PHDAM or PHIDAM HALDB partitions online, HALDB Online Reorganization (OLR) requires additional data sets to hold the reorganized data. The additional data sets are equal in number to the data sets being reorganized, excluding the ILDS.

In a PHDAM database, HALDB OLR increases the maximum number of data sets associated with a partition to twenty-one. In a PHIDAM database, which includes a primary index, HALDB OLR increases the maximum number of data sets associated with a partition to twenty-three. In either case, HALDB OLR only needs as many new data sets as exist in the partition at the time the reorganization process begins.

Each additional data set requires an additional DD name. To distinguish between the DD names for the data sets that HALDB OLR reorganizes and the DD names for the data sets into which HALDB OLR moves the reorganized data, HALDB OLR extends the HALDB naming convention to include the suffixes M through V for the DD names of the primary data sets and the suffix Y for the DD name for the additional primary index. The suffixes M through V and Y correspond in order to the standard HALDB DD name suffixes A through J and X.

Related Reading: For more information on HALDB OLR, see “HALDB Online Reorganization” on page 370.

HALDB Data Set Names

A HALDB partition uses a minimum of one, two, or three data sets and a maximum of one, eleven, or thirteen data sets depending on the type of HALDB—PSINDEX,

Establishing Naming Conventions

PHDAM, or PHIDAM—you are defining. The naming convention for HALDB data sets within a partition is designed to simplify the naming of multiple data sets.

DL/I pointers within the segment prefix that point into another partition use a halfword binary number as the target partition identification. DL/I must be able to correlate this number to the correct partition. By using a data set naming convention, DL/I can correlate the halfword binary number to the data set name for the partition. You specify the base name and the suffix is assigned by DL/I.

DL/I assigns the following suffixes:

- .ANNNNN for the first data set
- .BNNNNN for the second data set
- ...
- ...
- .JNNNNN for the tenth data set
- .XNNNNN for the primary index data set
- .LNNNNN for the ILDS

Where NNNNN is the five-digit numerical partition identifier assigned to each partition during partition definition. Note that this decimal number is not the processing sequence of the database. The maximum value of the partition identifier is 31 999.

Extended Naming Convention for Data Set When Using HALDB Online

Reorganization: To distinguish between the data sets that HALDB OLR reorganizes and the data sets into which HALDB OLR moves the reorganized data, HALDB OLR uses the same extended naming convention it uses for DD names: the characters M through V for the primary data sets and the character Y for the additional primary index. For data sets, IMS combines these characters with the partition ID to form the suffix that uniquely identifies each HALDB data set.

Related Reading: For more information about HALDB OLR naming conventions, see “Data Set Naming Conventions for HALDB Online Reorganization” on page 378.

Chapter 3. Review Process

One of the best ways to make sure a good database design is developed and effectively implemented is to review the design at various stages in its development. The sections in this chapter describe the reviews typically conducted during development of a database system. The types of reviews are:

- Design reviews 1, 2, 3, and 4
- Code inspections 1 and 2
- Security inspection
- Post-implementation review

In this Chapter:

- “The Design Review”
- “Design Review 1” on page 26
- “Design Review 2” on page 26
- “Design Review 3” on page 27
- “Design Review 4” on page 27
- “Code Inspection 1” on page 28
- “Who Attends Code Inspection 1” on page 28
- “Code Inspection 2” on page 28
- “Security Inspection” on page 29
- “Post-Implementation Review” on page 29

The Design Review

Design Reviews ensure that the functions being developed are adequate, the performance is acceptable, the installation standards met, and the project is understood and under control. Hold reviews during development of the initial database system and, afterward, whenever a program or set of programs is being developed to run against it.

Role of the Database Administrator

The role of database administration in the review process is an important one. Typically, a member of the database administration staff, someone not associated with the specific system being developed, moderates the reviews. The moderator does more than just conduct the meeting. The moderator also looks to see what impact development of this system has on existing or future systems. You, the database administrator responsible for developing the system, need to participate in all reviews.

Your role in the review process is to ensure that a good database design is developed and then effectively implemented. The role is ongoing and provides a supporting framework for the other database administration tasks described in this book.

General Information about Reviews

The sections of this chapter describe reviews typically held during system development. (For purposes of simplicity, “system” describes the object under review. In actuality, the “system” could be a program, set of programs, or an entire database system.) The number of reviews, who attends them, and their specific role

The Design Review

in the review will differ slightly from one installation to the next. What you need to understand is the importance of the reviews and the tasks performed at them. Here is some general information about reviews:

- People attending all reviews (in addition to database administrators) include a review team and the system designer. The review team generally has no responsibility for developing the system. The review team consists of a small group of people whose purpose is to ensure continuity and objectivity from one review to the next. The system designer writes the initial functional specifications.
- At the end of each review, make a list of issues raised during the review. These issues are generally change requirements. Assign each issue to a specific person for resolution, and set a target date for resolution. If certain issues require major changes to the system, schedule other reviews until you resolve all major issues.
- If you have a data dictionary, update it at the end of each review to reflect any decisions that you made. The dictionary is an important aid in keeping information current and available especially during the first four reviews when you make design decisions.

Design Review 1

The first design review takes place after initial functional specifications for the system are complete. Its purpose is to ensure that all user requirements have been identified and that design assumptions are consistent with objectives. No detailed design for the system is or should be available at this point. The review of the specifications will determine whether the project is ready to proceed to a more detailed design. When design review 1 concludes successfully, its output is an approved set of initial *functional* specifications.

People who attend design review 1, in addition to the regular attendees, include someone from the organization that developed the requirement and anyone participating in the development of detailed design. You are at the review primarily for information. You also look at:

- The relationship between data elements
- Whether any of the needed data already exists

Design Review 2

The second design review takes place after final *functional* specifications for the system are complete. This means the overall logic for each program in the system is defined, as well as the interface and interactions between programs. Audit and security requirements are defined at this point, along with most data requirements. When design review 2 is successfully concluded, its output is an approved set of final functional specifications.

Everyone who attended design review 1 should attend design review 2. People from test and maintenance groups attend as observers to begin getting information for test case design and maintenance. Those concerned with auditing and security can also attend.

Your role in this review is still primarily to gather information. You also look at:

- Whether the specifications meet user requirements
- Whether the relationship between data items is correct
- Whether any of the required data already exists
- Whether audit and security requirements are consistent with user requirements

- Whether audit and security requirements can be implemented

Design Review 3

The third design review takes place after initial *logic* specifications for the system are complete. At this point, high level pseudo code or flowcharts are complete. These can only be considered complete when major decision points in the logic are defined, calls or references to external data and modules are defined, and the general logic flow is known. All modules and external interfaces are defined at this point, definition of data requirements is complete, and database and data files are designed. Initial test and recovery plans are available; however, no code has been written. When design review 3 concludes successfully, its output is an approved set of initial logic specifications.

Everyone who attended design review 2 should attend design review 3. If the project is large, those developing detailed design need only be present during the review of their portion of the project.

It is possible now that logic specifications are available.

Your role in this review is to ensure that the flow of transactions is consistent with the database design you are creating.

At this point in the design review process, you are designing hierarchies and starting to design the database. These tasks are described in Chapter 5, “Analyzing Data Requirements,” on page 49, Chapter 6, “Choosing Full-Function Database Types,” on page 59, Chapter 8, “Choosing Optional Database Functions,” on page 157, and Chapter 9, “Designing Full-Function Databases,” on page 245.

Design Review 4

The fourth design review takes place after design review 3 is completed and all interested parties are satisfied that system design is essentially complete. No special document is examined at this review, although final functional specifications and either initial or final logic specifications are available. The primary objective of this review is to make sure that system performance will be acceptable.

At this point in the development process, sufficient flexibility exists to make necessary adjustments to the design, since no code exists but detailed design is complete. Although some design changes undoubtedly occur once coding is begun; these changes should not impact the entire system. Although no code exists at this point, you can and should run tests to check that the database you have designed will produce the results you expect.

When design review 4 concludes successfully, database design is considered complete.

The people who attend all design reviews (moderator, review team, database administrator, and system designer) should attend design review 4. Others attend only as specific detail is required.

At this point in the review process, you are almost finished with the database administration tasks along with designing and testing your database. These tasks are described in Chapter 5, “Analyzing Data Requirements,” on page 49, Chapter 6, “Choosing Full-Function Database Types,” on page 59, and Chapter 12, “Developing Test Databases,” on page 311.

Code Inspection 1

The first code inspection takes place after final logic specifications for the system are complete.

At this point, no code is written but the final functional specifications have been interpreted. Both pseudo code and flowcharts have a statement or logic box for every 5 to 25 lines of assembler language code, 5 to 15 lines of COBOL code, or 5 to 15 lines of PL/I code that needs writing. In addition, module prologues are written, and entry and exit logic along with all data areas are defined.

The objective of this review is to ensure that the correctly developed logic interprets the functional specification. Code inspection 1 also provides an opportunity to review the logic flow for any performance implications or problems. When code inspection 1 successfully concludes, its output is an approved set of final logic specifications.

Who Attends Code Inspection 1

Code inspection 1 is attended primarily by those doing the coding. People who attend all design reviews (moderator, review team, database administrator, and system designer) also attend the code inspection 1. Testing people present the test cases that will be used to validate the code, while maintenance people are there to learn and evaluate maintainability of the database.

Your role in this review is now a less active one than it has been. You are there to ensure that everyone adheres to the use of data and access sequences defined in the previous reviews.

At this point in the review process, you are starting the database administration tasks defined in Chapter 12, “Developing Test Databases,” on page 311, Chapter 11, “Implementing Database Design,” on page 295, and Chapter 13, “Loading Databases,” on page 315.

Code Inspection 2

The code inspection 2 takes place after coding is complete and before testing by the test organization begins. The objective of the second code inspection is to make sure module logic matches pseudo code or flowcharts. Interface and register conventions along with the general quality of the code are checked. Documentation and maintainability of the code are evaluated.

Everyone who attended code inspection 1 should attend code inspection 2.

Your role in this review is the same as your role in code inspection 1.

At this point in the review process, you are almost finished with the database administration tasks of developing a test database, implementing the database design, and loading the database.

During your testing of the database, you should run the DB monitor (described in Chapter 14, “Monitoring Databases,” on page 341) to make sure your database still meets the performance expectations you have established.

Security Inspection

The security inspection is optional but highly recommended if security is a significant concern. Security inspections can take place at any appropriate point in the system development process. Define security strategy early, and check its implementation during design reviews. This particular security inspection takes place after all unit and integration testing is complete. The purpose of the review is to look for any code that violates the security of system interfaces, secured databases, tables, or other high-risk items.

People who attend the security inspection review include the moderator, system designer, designated security officer, and database administrator. Because the database administrator is responsible for implementing and monitoring the security of the database, you might, in fact, be the designated security officer. If security is a significant concern, you might prefer that the review team not attend this inspection.

During this and other security inspection, you are involved in the database administration task of establishing security defined in Chapter 4, “Security,” on page 31.

Post-Implementation Review

It is highly recommended that you conduct a post-implementation review. The post-implementation review is typically held about six months after the database system is running. Its objective is to make sure the system is meeting user requirements.

Everyone who has been involved in design and implementation of the database system should attend the post-implementation review. If the system is not meeting user requirements, the output of this review should be a plan to correct design or performance problems to meet user requirements.

Chapter 4. Security

The two aspects of database security are as follows:

- *User verification* (how you establish that the person using an online database is in fact the person you have authorized)
- *User authority* (once you have verified the user's identity, how you control what is seen—and what can be done with what is seen)

This chapter deals primarily with how you can control a user's view of data and the user's actions with respect to the data.

This chapter examines the following areas:

- "Restricting the Scope of Data Access"
- "Restricting Processing Authority"
- "Restricting Access by Non-IMS Programs" on page 33
- "Using the Dictionary to Help Establish Security" on page 33

Related Reading: If you use CICS, see *CICS RACF® Security Guide* for information on establishing security.

Restricting the Scope of Data Access

The PCB defines a program's (and therefore the user's) view of the database. The PCB can be thought of as a "mask" over the data structure defined by the DBD, hiding certain parts of it. Therefore, it is possible, simply by limiting the scope of the PCB, to limit the user's access to (and even knowledge of) elements of the database you need to restrict.

Figure 14 on page 32 shows an example. The top of the figure shows the hierarchical structure for a PAYROLL database as seen by you and defined by the DBD. For certain applications, it is not necessary (nor desirable) to access the SALARY segment. By omitting SENSEG statement in the DB PCB for the SALARY segment, you can make it seem that this segment simply does not exist. By doing this, you have denied unauthorized users access to the segment, and you have denied users knowledge of its very existence.

For this method to be successful, the segment being masked off must not be in the search path of an accessed segment. If it is, then the application is made aware of at least the key of the segment to be "hidden."

With field-level sensitivity, you can achieve the same masking effect at the *field* level. If SALARY and NAME were in the same segment, you could still restrict access to the SALARY field without denying access to other fields in the segment.

Restricting Processing Authority

After you have controlled the scope of data a user has access to, you can also control authority within that scope. Controlling authority allows you to decide what processing actions against the data a given user is permitted. For example, you could give some application programs authority only to read segments in a database, while you give others authority to update or delete segments. You can do this through the PROCOPT parameter of the SENSEG statement and through the

Restricting Processing Authority

PCB statement. The PROCOPT statement tells IMS what actions you will permit against the database. A program can do what is declared in the PROCOPT.

In addition to restricting access and authority, the number of sensitive segments and the processing option specified can have an impact on data availability. To achieve maximum data availability, the PSB should be sensitive only to the segments required and the processing option should be as restrictive as possible.

For example, the DBD in Figure 13 describes a payroll database that stores the name, address, position, and salary of employees. The hierarchical structure of the database record is shown in Figure 14.

```
DBD  NAME=PAYROLL,...  
DATASET ...  
SEGM  NAME=NAME,PARENT=0...  
FIELD  NAME=  
SEGM  NAME=ADDRESS,PARENT=NAME,...  
FIELD  NAME=  
SEGM  NAME=POSITION,PARENT=NAME,...  
FIELD  NAME=  
SEGM  NAME=SALARY,PARENT=NAME,...  
FIELD  NAME=  
:  
:
```

Figure 13. DBD for Payroll Database




Figure 14. Payroll Database Record without a Mask

If an application needs access to the name, address, and position of employees, but not the salary, use the SENSEG statement of the DB PCB to make the application sensitive to only the name, address, and position segments. The SENSEG statements on the DB PCB creates a mask over the database record hiding segments from application. Figure 15 shows the DB PCB that masks the SALARY segment of the payroll database from the application.

```
PCB  TYPE=DB,DBDNAME=PAYROLL,...  
SENSEG  NAME=NAME,PARENT=0,...  
SENSEG  NAME=ADDRESS,PARENT=NAME,...  
SENSEG  NAME=POSITION,PARENT=NAME,...  
:  
:
```

Figure 15. PCB for Payroll Database

Figure 16 on page 33 shows what the payroll database record looks like to the application based on the DB PCB. It looks just like the database record in Figure 14 except that the SALARY segment is hidden.

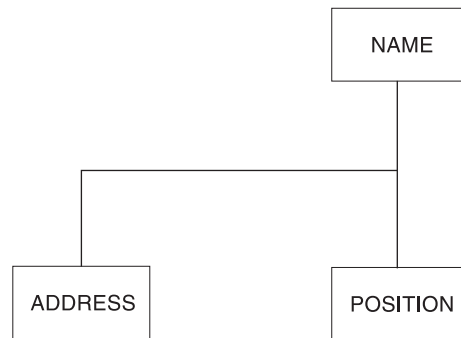


Figure 16. Payroll Database Record with SALARY Segment Masked

Restricting Access by Non-IMS Programs

One potential security exposure is from people attempting to access IMS data sets with non-IMS programs. Two methods of protecting against this exposure are data set password protection and database encryption.

Protecting Data with VSAM Passwords

You can take advantage of VSAM password protection to prevent non-IMS programs from reading VSAM data sets on which you have your IMS databases. To protect data with VSAM passwords, specify password protection for your VSAM data sets and code `PASSWD=YES` on the DBD statement. IMS then passes the DBD name as the password. If you specify `PASSWD=NO` on the DBD statement, the console operator is prompted to provide a password to VSAM each time the data set is opened.

This method is only useful in the batch environment, and VSAM password checking is bypassed entirely in the online system. (If you have RACF installed, you can use it to protect VSAM data sets.)

Details of the `PASSWD` parameter of the DBD statement can be found in *IMS Version 9: Utilities Reference: System*.

Encrypting Your Database

Another precaution you can take against non-IMS programs reading DL/I databases is to encrypt the databases. You can encrypt DL/I segments using your own encryption routine, entered at the segment edit/compression exit. Before segments are written on the database, IMS passes control to your routine, which encrypts them. Then, each time they are retrieved, they are decrypted by your routine before presentation to the application program.

Do not change the key or the location of the key field in index databases or in root segments of HISAM data bases.

You can learn more about segment edit/compression routines in “Segment Edit/Compression Exit Routine” on page 218.

Using the Dictionary to Help Establish Security

The dictionary monitors relationships among entities in your computing environment (such as, which programs use which data elements). This ability makes the dictionary the ideal tool to administer security.

Using the Dictionary to Help Establish Security

You can use the dictionary to define your authorization matrixes. Through the extensibility feature, you can define terminals, programs, users, data, and their relationships to each other. In this way, you can produce reports that show: dangerous trends, who uses what from which terminal, and which user gets what data. For each user, the dictionary could be used to list the following information:

- Programs that can be used
- Types of transactions that can be entered
- Data sets that can be read
- Data sets that can be modified
- Categories of data within a data set that can be read
- Categories of data that can be modified

Part 2. Administering IMS Databases

| | |
|---|--------|
| Chapter 5. Analyzing Data Requirements | 49 |
| Local View | 49 |
| Local View 1. Current Roster. | 50 |
| List of Current Roster Data Elements. | 50 |
| Current Roster Mappings | 51 |
| Local View 2. Schedule of Classes | 51 |
| List of Schedule of Classes Data Elements | 51 |
| Schedule of Classes Mappings | 52 |
| Local View 3. Instructor Skills Report. | 52 |
| List of Instructor Skills Report Data Elements. | 52 |
| Instructor Skills Report Mappings | 52 |
| Local View 4. Instructor Schedules | 53 |
| List of Instructor Schedules Data Elements | 53 |
| Instructor Schedules Mappings | 53 |
| Designing a Conceptual Data Structure | 53 |
| Implementing the Structure with DL/I | 55 |
| Assigning Data Elements to Segments | 55 |
| Resolving Data Conflicts | 56 |
| Analyzing Requirements for Secondary Indexes. | 56 |
| Analyzing Requirements for Logical Relationships | 56 |
| Chapter 6. Choosing Full-Function Database Types | 59 |
| Sequential Storage Method | 60 |
| Direct Storage Method | 60 |
| Databases Supported with DBCTL. | 60 |
| Databases Supported with DCCTL | 61 |
| Performance Considerations Overview | 61 |
| HSAM Databases | 64 |
| When to Use HSAM | 65 |
| How an HSAM Record Is Stored | 65 |
| DL/I Calls against an HSAM Database | 67 |
| No Sequence Field Defined | 68 |
| Sequence Field Defined | 68 |
| HISAM Databases | 69 |
| When to Use HISAM. | 69 |
| How a HISAM Record is Stored | 69 |
| Accessing Segments. | 72 |
| Inserting Root Segments Using VSAM | 72 |
| A Free Logical Record Exists. | 72 |
| No Free Logical Record Exists | 73 |
| Inserting Dependent Segments | 74 |
| Deleting Segments | 76 |
| Replacing Segments. | 78 |
| Criteria for Selecting HISAM | 78 |
| SHSAM, SHISAM and GSAM Databases | 78 |
| Situation 1 - Converting from a non-database system to IMS | 78 |
| Situation 2 - Passing data | 79 |
| SHSAM Databases | 79 |
| SHISAM Databases | 79 |
| SHISAM IMS Symbolic Checkpoint Call. | 80 |
| GSAM Databases. | 80 |
| GSAM IMS Symbolic Checkpoint Call | 80 |
| HDAM, PHDAM, HIDAM, and PHIDAM Databases. | 82 |

| | |
|---|------------|
| Maximum Sizes of HD Databases | 83 |
| DL/I Calls Issuable Against HD Databases | 84 |
| When to Use HDAM and PHDAM | 84 |
| When to Use HIDAM and PHIDAM | 85 |
| What You Need to Know About HD Databases | 85 |
| Types of Pointers You Can Specify | 85 |
| Hierarchic Forward Pointers | 86 |
| Hierarchic Forward and Backward Pointers | 87 |
| Physical Child First Pointers | 88 |
| Physical Child First and Last Pointers | 89 |
| Physical Twin Forward Pointers | 91 |
| Physical Twin Forward and Backward Pointers | 92 |
| Mixing Pointers | 93 |
| Sequence of Pointers in a Segment's Prefix | 94 |
| General Format of HD Databases and Use of Special Fields | 95 |
| How HDAM and PHDAM Records Are Stored | 98 |
| When Not Enough Root Storage Room Exists | 100 |
| How HIDAM and PHIDAM Records Are Stored | 100 |
| Loading a HIDAM or PHIDAM Database | 101 |
| Creating an Index Segment | 102 |
| Use of RAPs in a HIDAM Database | 102 |
| Accessing Segments | 103 |
| Qualified Calls | 103 |
| Unqualified Calls | 104 |
| Inserting Root Segments | 104 |
| Inserting Root Segments into an HDAM or PHDAM Database | 104 |
| Inserting Root Segments Into a HIDAM or PHIDAM Database | 104 |
| Updating the Space Management Fields When a Root Segment Is Inserted | 105 |
| Inserting Dependent Segments | 106 |
| Deleting Segments | 107 |
| Replacing Segments | 107 |
| How the HD Space Search Algorithm Works | 107 |
| Root Segment | 108 |
| Dependent Segment | 108 |
| Second-Most Desirable Block | 108 |
| Locking Protocols | 109 |
| Locking to Provide Program Isolation | 109 |
| Locking for Q Command Codes | 110 |
| Resource Locking Considerations with Block Level Sharing | 110 |
| Data Sharing Impact on Locking | 111 |
| Locking in HDAM, PHDAM, HIDAM, and PHIDAM Databases | 111 |
| Locking for Secondary Indexes | 111 |
| Managing I/O Errors | 111 |
| Chapter 7. Choosing Fast Path Database Types | 113 |
| Data Entry Databases | 113 |
| DEDB Functions | 114 |
| DEDB Areas | 114 |
| Areas and the DEDB Format | 114 |
| Opening and Preopening DEDB Areas | 115 |
| Stopping and Starting DEDBs and DEDB Areas | 116 |
| Restarting and Reopening Areas After an IRLM Failure | 116 |
| Read and Write Errors in DEDB Areas | 117 |
| Record Deactivation | 118 |
| Non-Recovery Option | 118 |

| | |
|---|-----|
| Area Data Set Replication | 119 |
| DEDBs and Data Sharing | 119 |
| Fixed- and Variable-Length Segments in DEDBs | 120 |
| Examples of Defining Segments | 120 |
| Parts of a DEDB Area | 121 |
| Root Addressable Part | 123 |
| Independent Overflow Part | 123 |
| Sequential Dependent Part | 123 |
| CI and Segment Formats | 123 |
| Root Segment Storage | 126 |
| Direct Dependent Segment Storage | 126 |
| Sequential Dependent Segment Storage | 127 |
| Enqueue Level of Segment CIs | 127 |
| DEDB Space Search Algorithm | 129 |
| DEDB Insert Algorithm | 129 |
| DEDB Free Space Algorithm | 130 |
| Reorganization | 130 |
| Segment Deletion | 131 |
| Managing Unusable Space with IMS Tools | 131 |
| DL/I Calls against a DEDB | 131 |
| Mixed Mode Processing | 131 |
| Main Storage Databases (MSDBs) | 132 |
| When to Use an MSDB | 133 |
| MSDBs Storage | 133 |
| MSDB Record Storage | 134 |
| Saving MSDBs for Restart | 135 |
| DL/I Calls against an MSDB | 135 |
| Rules for Using an SSA | 135 |
| Insertion and Deletion of Segments | 136 |
| Combination of Binary and Direct Access Methods | 136 |
| Position in an MSDB | 137 |
| The Field Call | 138 |
| Call Sequence Results | 138 |
| Fast Path Virtual Storage Option | 139 |
| Restrictions Using VSO DEDB Areas | 139 |
| Defining a VSO DEDB Area | 140 |
| VSO DEDB Areas and the PREOPEN and NOPREO Keywords | 141 |
| Sharing of VSO DEDB Areas | 142 |
| The Coupling Facility and Shared Storage | 142 |
| Duplexing Structures | 143 |
| Automatic Altering of Structure Size | 143 |
| System-Managed Rebuild | 144 |
| Private Buffer Pools | 144 |
| Defining a VSO Cache Structure Name | 144 |
| Coupling Facility Structure Naming Convention | 144 |
| Examples of Defining Coupling Facility Structures | 144 |
| Registering a Cache Structure Name with DBRC | 145 |
| Defining a Private Buffer Pool Using the DFSVSMxx IMS.PROCLIB Member | 146 |
| Defining a Private Buffer Pool for a Multi-Area Structure | 147 |
| Acquiring and Accessing Data Spaces for VSO DEDB Areas | 147 |
| Acquiring a Data Space | 147 |
| Accessing a Data Space | 148 |
| Resource Control and Locking | 148 |
| Preopen Areas and VSO Areas in a Data Sharing Environment | 149 |
| Input/Output Processing With VSO | 150 |

| | |
|--|------------|
| Input Processing | 150 |
| Output Processing | 150 |
| The PRELOAD Option | 151 |
| I/O Error Processing | 151 |
| Checkpoint Processing | 152 |
| VSO Options Across IMS Restart. | 152 |
| Emergency Restart Processing | 152 |
| VSO Options with XRF | 153 |
| Fast Path Synchronization Points. | 153 |
| Phase 1 - Build Log Record. | 153 |
| Phase 2 - Write Record to System Log | 153 |
| Managing I/O Errors and Long Wait Times | 154 |
| Registering Fast Path Databases in DBRC | 155 |
| Chapter 8. Choosing Optional Database Functions | 157 |
| Logical Relationships | 157 |
| Logical Relationship Types | 158 |
| Unidirectional Logical Relationships | 159 |
| Bidirectional Physically Paired Logical Relationship | 160 |
| Bidirectional Virtually Paired Logical Relationship | 161 |
| Logical Relationship Pointer Types | 162 |
| Logical Parent Pointer. | 162 |
| Logical Child Pointer | 164 |
| Physical Parent Pointer | 165 |
| Logical Twin Pointer | 166 |
| Indirect Pointers | 167 |
| Paths in Logical Relationships | 168 |
| The Logical Child Segment | 169 |
| Segment Prefix Information for Logical Relationships | 170 |
| Sequence of Pointers in a Segment's Prefix. | 170 |
| Counter Used in Logical Relationships. | 170 |
| Intersection Data. | 170 |
| Fixed Intersection Data | 171 |
| Variable Intersection Data | 171 |
| FID, VID, and Physical Pairing. | 172 |
| Recursive Structures: Same Database Logical Relationships | 172 |
| Defining Sequence Fields for Logical Relationships | 176 |
| Logical Parent Sequence Fields | 176 |
| Real Logical Children Sequence Fields | 177 |
| Virtual Logical Children Sequence Fields | 177 |
| Control Blocks for Logical Relationships | 177 |
| Specifying Logical Relationships in the Physical DBD | 178 |
| Specifying Bidirectional Logical Relationships | 181 |
| Checklist of Rules for Defining Logical Relationships in Physical Databases | 181 |
| Specifying Logical Relationships in the Logical DBD. | 182 |
| Checklist of Rules for Defining Logical Databases | 183 |
| Choosing Replace, Insert, and Delete Rules for Logical Relationships | 187 |
| Performance Considerations for Logical Relationships | 189 |
| Logical Parent Pointers | 189 |
| KEY/DATA Considerations | 190 |
| Sequencing Logical Twin Chains | 191 |
| Placement of the Real Logical Child in a Virtually Paired Relationship | 192 |
| Secondary Indexes | 192 |
| Why Secondary Indexes? | 192 |
| Characteristics of Secondary Indexes | 194 |

| | |
|--|-----|
| Segments Used for Secondary Indexes | 194 |
| How the Hierarchy Is Restructured | 197 |
| Secondary Data Structure | 198 |
| Secondary Processing Sequence. | 198 |
| How a Secondary Index Is Stored | 198 |
| Format and Use of Fields in a Pointer Segment | 199 |
| Delete Byte. | 200 |
| Pointer Field | 200 |
| Constant Field. | 200 |
| Search Field | 200 |
| Subsequence Field | 200 |
| Duplicate Data Field | 201 |
| Concatenated Key Field | 201 |
| User Data in Pointer Segments | 202 |
| Making Keys Unique Using System Related Fields | 202 |
| Using the /SX Operand | 202 |
| Using the /CK Operand | 202 |
| Suppressing Index Entries: Sparse Indexing. | 204 |
| Advantages of Sparse Indexing | 204 |
| How to Specify a Sparse Index | 204 |
| How the Secondary Index Is Maintained | 205 |
| Inserting a Source Segment | 205 |
| Deleting a Source Segment. | 205 |
| Replacing a Source Segment | 205 |
| Processing a Secondary Index as a Separate Database | 206 |
| Sharing Secondary Index Databases | 207 |
| Using the INDICES= Parameter | 207 |
| Compare Process and Performance | 208 |
| Using Secondary Indexes with Logical Relationships | 209 |
| Using Secondary Indexes with Variable-Length Segments | 210 |
| Considerations When Using Secondary Indexing | 210 |
| How to Specify Use of Secondary Indexing in the DBD. | 211 |
| DBD for the EDUC Database | 211 |
| DBD for the SINDX Database | 213 |
| Choosing Secondary Indexes Versus Logical Relationships | 214 |
| When to Use a Secondary Index | 214 |
| When to Use a Logical Relationship | 214 |
| Variable-Length Segments | 215 |
| How to Specify Variable-Length Segments | 215 |
| How Variable-Length Segments Are Stored and Processed | 216 |
| When to Use Variable-Length Segments | 217 |
| What Application Programmers Need to Know about Variable-Length Segments | 217 |
| Adding or Converting to Variable-Length Segments | 218 |
| Segment Edit/Compression Exit Routine | 218 |
| Things to Consider Before Using the Segment Edit/Compression Exit Routine | 219 |
| Preventing Split Segments from Impacting Performance | 220 |
| How to Specify the Segment Edit/Compression Exit Routine. | 220 |
| Converting to the Segment Edit/Compression Exit Routine | 220 |
| Data Capture Exit Routines | 221 |
| DBD Parameters for Data Capture Exit Routines | 221 |
| The EXIT= Parameter | 222 |
| The VERSION= Parameter | 222 |
| Call Sequence of Data Capture Exit Routines | 223 |
| Data Passed To and Captured By the Data Capture Exit Routine | 224 |

| | |
|---|-----|
| Data Capture Call Functions | 224 |
| Cascade Delete When Crossing Logical Relationships | 225 |
| Data Capture Exit Routines and Logically Related Databases | 225 |
| Converting to Data Capture Exit Routines | 225 |
| Field-Level Sensitivity | 226 |
| Using Field-Level Sensitivity as a Mapping Interface. | 226 |
| Using Field-Level Sensitivity with Variable-Length Segments. | 226 |
| How to Specify Use of Field-Level Sensitivity in the DBD and PSB | 227 |
| Retrieving Segments Using Field-Level Sensitivity | 228 |
| Replacing Segments Using Field-Level Sensitivity | 228 |
| Inserting Segments Using Field-Level Sensitivity | 229 |
| Using Field-Level Sensitivity When Fields Overlap | 230 |
| Using Field-Level Sensitivity When Path Calls Are Issued. | 230 |
| Using Field-Level Sensitivity with Logical Relationships | 230 |
| Using Field-Level Sensitivity with Variable-Length Segments. | 231 |
| Retrieving Missing Fields. | 231 |
| Replacing Missing Fields. | 232 |
| Inserting Missing Fields | 233 |
| Retrieving Partially Present Fields | 234 |
| Replacing Partially Present Fields | 234 |
| General Considerations for Using Field-Level Sensitivity | 235 |
| Multiple Data Set Groups | 235 |
| Why Use Multiple Data Set Groups? | 236 |
| HD Databases Using Multiple Data Set Groups | 237 |
| How HD Records Are Stored in Multiple Data Set Groups | 238 |
| Specifying Use of Multiple Data Set Groups in HD and PHD Databases | 239 |
| VSAM KSDS CI Reclaim For Full-Function Databases | 241 |
| HALDB Single Partition Processing | 242 |
| Logical Relationships in Single Partition Processing | 242 |
| Secondary Indexes in Single Partition Processing | 242 |
| Partition Selection | 243 |
| Examples of Single Partition Processing | 243 |
| Integrated HALDB Online Reorganization Function | 244 |
| Storing XML Data in IMS Databases | 244 |
| Chapter 9. Designing Full-Function Databases. | 245 |
| Specifying Free Space (HDAM, PHDAM, HIDAM, and PHIDAM Only) | 245 |
| Estimating the Size of the Root Addressable Area (HDAM or PHDAM Only) | 246 |
| Determining Which Randomizing Module to Use (HDAM and PHDAM Only) | 247 |
| Write Your Own Randomizing Module | 247 |
| Assess the Effectiveness of the Randomizing Module | 247 |
| Choosing HDAM or PHDAM Options | 248 |
| Minimizing I/O Operations | 248 |
| Maximizing Packing Density | 248 |
| Choosing a Logical Record Length for a HISAM Database | 249 |
| Logical Record Length Considerations. | 249 |
| Rules to Observe | 251 |
| Calculating How Many Logical Records Are Needed to Hold a Database | 252 |
| Record | 252 |
| Specifying Logical Record Length | 252 |
| Choosing a Logical Record Length for HD Databases | 252 |
| Determining the Size of CIs and Blocks | 252 |
| Buffering Options | 253 |
| Multiple Buffers in Virtual Storage | 253 |
| "Use" Chain | 253 |
| The Buffer Handler | 253 |

| | |
|---|------------|
| Background Write Option | 254 |
| Shared Resource Pools | 254 |
| Using Separate Subpools | 254 |
| Hiperspace Buffering | 254 |
| Buffer Size | 254 |
| Buffer Numbers | 255 |
| VSAM Buffer Sizes | 255 |
| OSAM Buffer Sizes | 256 |
| Specifying Buffers | 256 |
| OSAM Sequential Buffering | 257 |
| Sequential Buffering Introduction | 257 |
| Benefits of Sequential Buffering | 258 |
| Programs That Can Benefit from SB | 258 |
| Typical Productivity Benefits of SB | 258 |
| Flexibility of SB Use | 259 |
| How SB Buffers Data | 259 |
| What SB Buffers | 259 |
| Conditional Activation and Periodical Evaluation of SB | 259 |
| Role of the SB Buffer Handler | 260 |
| Virtual Storage Considerations for SB | 260 |
| How to Request the Use of SB | 261 |
| Requesting SB During PSB Generation | 261 |
| Requesting SB With SB Control Statements | 262 |
| Requesting SB with an SB Initialization Exit Routine | 262 |
| SB Options or Parameters Provided by Several Sources | 263 |
| Using SB in an Online System | 263 |
| Disallowing the Use of SB | 263 |
| VSAM Options | 264 |
| Optional Functions Specified in the OPTIONS Control Statement | 264 |
| Using Background Write (BGWRT Parameter) | 264 |
| Choosing an Insert Strategy (INSERT Parameter) | 265 |
| Using the IMS Trace Parameters | 266 |
| Determining Which Dump Option to Use (DUMP Parameter) | 266 |
| Deciding Whether to Fix VSAM Database Buffers and IOBs in Storage (VSAMFIX Parameter) | 266 |
| Using Local Shared Resources (VSAMPLS Parameter) | 266 |
| Optional Functions Specified in the POOLID, DBD, and VSRBF Control Statements | 266 |
| Optional Functions Specified in the Access Method Services DEFINE CLUSTER Command | 267 |
| Specifying that 'Fuzzy' Image Copies Can be Taken with the Database Image Copy 2 (DFSUDMT0) | 267 |
| Specifying Free Space for a KSDS (FREESPACE Parameter) | 267 |
| Specifying Whether Data Set Space Is Pre-formatted for Initial Load (SPEED RECOVERY Parameter) | 267 |
| Specifying Whether Index Set Records Are Replicated | 268 |
| OSAM Options | 269 |
| Dump Option (DUMP Parameter) | 269 |
| Deciding Which FIELD Statements to Code in the DBD | 269 |
| Planning for Maintenance | 269 |
| Chapter 10. Designing Fast Path Databases | 271 |
| Designing a Data Entry Database (DEDB) | 271 |
| DEDB Design Guidelines | 271 |
| DEDB Area Design Guidelines | 272 |
| Determining the Size of the CI | 273 |

| | |
|--|------------|
| Determining the Size of the UOW | 274 |
| SDEP CI Preallocation and Reporting | 274 |
| Processing Option P (PROCOPT=P) | 275 |
| DEDB Randomizing Routine Design | 275 |
| Multiple Copies of an Area Data Set | 276 |
| Record Deactivation | 277 |
| Physical Child Last Pointers | 277 |
| Subset Pointers | 277 |
| Designing a Main Storage Database (MSDB) | 277 |
| Calculating Virtual Storage Requirements for an MSDB | 278 |
| Calculating Buffer Requirements | 278 |
| Calculating the Storage for an Application I/O Area | 278 |
| Understanding Resource Allocation, a Key to Performance | 279 |
| Designing to Minimize Resource Contention. | 280 |
| Choosing MSDBs to Load and Page-Fix | 281 |
| Auxiliary Storage Requirements for an MSDB | 283 |
| High-Speed Sequential Processing (HSSP) | 283 |
| Why HSSP? | 284 |
| Limitations and Restrictions When Using HSSP | 284 |
| Using HSSP | 285 |
| HSSP Processing Option H (PROCOPT=H). | 285 |
| Image-Copy Option | 285 |
| UOW Locking | 286 |
| Private Buffer Pools | 286 |
| Designing a DEDB or MSDB Buffer Pool | 286 |
| Buffer Requirements | 287 |
| Normal Buffer Allocation (NBA) | 287 |
| Overflow Buffer Allocation (OBA) | 287 |
| Fast Path Buffer Allocation Algorithm | 287 |
| System Buffer Allocation (DBFX) | 288 |
| Determining the Fast Path Buffer Pool Size | 288 |
| Fast Path Buffer Performance Considerations | 288 |
| The NBA Limit and Sync Point. | 289 |
| The DBFX Value and the Low Activity Environment | 289 |
| Designing a DEDB Buffer Pool in the DBCTL Environment | 290 |
| Buffer Requirements in a DBCTL Environment. | 290 |
| Normal Buffer Allocation for BMPs | 290 |
| Normal Buffer Allocation for CCTL Regions and Threads | 290 |
| Overflow Buffer Allocation for BMPs. | 291 |
| Overflow Buffer Allocation for CCTL Threads | 291 |
| Fast Path Buffer Allocation Algorithm for BMPs | 291 |
| Fast Path Buffer Allocation Algorithm for CCTL Threads | 292 |
| System Buffer Allocation (SBA) | 292 |
| Determining the Size of the Fast Path Buffer Pool for DBCTL | 292 |
| Fast Path Buffer Performance Considerations for DBCTL | 293 |
| The NBA/FPB Limit and Sync Point in a DBCTL Environment | 293 |
| Low Activity and the DBFX Value in a DBCTL Environment | 293 |
| A Note on Fast Path Buffer Allocation in IMS Regions | 294 |
| Chapter 11. Implementing Database Design | 295 |
| Coding Database Descriptions as Input for the DBDGEN Utility | 295 |
| The DBD Statement | 296 |
| The DATASET Statement | 296 |
| The SEGM Statement. | 297 |
| The FIELD Statement | 297 |
| The LCHILD Statement | 297 |

| | | |
|---|--|------------|
| | The XDFLD Statement | 298 |
| | The DBDGEN and END Statements | 298 |
| I | Implementing HALDB Design | 298 |
| I | Creating HALDBs with the HALDB Partition Definition Utility | 298 |
| I | Creating HALDB Partitions With the Partition Definition Utility | 299 |
| | Adding HALDB Partitions to an Existing HALDB | 302 |
| | Finding, Viewing, Sorting, Copying, Modifying, Deleting, and Printing HALDB Partitions. | 302 |
| | Defining and Modifying Data Set Groups | 303 |
| | Number of Data Sets in a HALDB Partition | 303 |
| | Exporting Database Definitions | 303 |
| | Importing Database Definitions | 303 |
| | Viewing IMS DDNAME Concatenation | 303 |
| I | Choosing IMS RECON and DBDLIB Libraries | 304 |
| | Deleting Database Information. | 304 |
| | Allocating an ILDS | 304 |
| | Coding Program Specification Blocks as Input to the PSBGEN Utility | 305 |
| | The Alternate PCB Statement | 306 |
| | The Database PCB Statement. | 306 |
| | The SENSEG Statement. | 307 |
| | The SENFLD Statement | 307 |
| | The PSBGEN Statement. | 307 |
| | The END Statement | 307 |
| | Building the Application Control Blocks (ACBGEN) | 308 |
| | Defining Generated Program Specification Blocks for SQL Applications. | 309 |
| | Chapter 12. Developing Test Databases | 311 |
| | Test Requirements | 311 |
| | What Kind of Database? | 312 |
| | What Kind of Sample Data? | 312 |
| | What Kind of Application Program? | 312 |
| | Designing, Creating, and Loading a Test Database | 312 |
| | Using Testing Standards | 312 |
| | Using IBM Programs to Develop a Test Database | 313 |
| | Using the Data Extraction, Processing, and Restructuring System. | 313 |
| I | Using IBM VisualAge Generator | 313 |
| I | IMS Application Development Facility II | 313 |
| I | File Manager for z/OS for IMS Data. | 314 |
| | Using the DL/I Test Program, DFSDDLTO. | 314 |
| | Chapter 13. Loading Databases | 315 |
| | Estimating the Minimum Size of the Database | 315 |
| | Step 1. Calculate the Size of an Average Database Record | 315 |
| | Determining Segment Size | 315 |
| | Determining Segment Frequency. | 316 |
| | Determining Average Database Record Size | 317 |
| | Step 2. Determine Overhead Needed for CI Resources | 317 |
| | Step 3. Determine the Number of CIs or Blocks Needed | 318 |
| | HISAM: Determining the Number of CIs or Blocks Needed | 318 |
| | HIDAM or PHIDAM: Determining the Number of CIs or Blocks Needed | 319 |
| | HIDAM or PHIDAM Index: Calculating the Space Needed | 319 |
| | HDAM or PHDAM: Determining the Amount of Space Needed | 319 |
| | Secondary Index: Determining the Amount of Space Needed | 321 |
| | Step 4. Determine the Number of Blocks or CIs Needed for Free Space | 321 |
| | Step 5. Determine the Amount of Space Needed for Bit Maps | 321 |
| | Allocating Data Sets | 322 |

| | |
|--|------------|
| Allocating OSAM Data Sets | 322 |
| Example of Allocating Multi-Volume OSAM Data Sets on Non-SMS Managed DASD | 323 |
| Example of Allocating Multi-Volume OSAM Data Sets on SMS-Managed DASD | 324 |
| Cautions When Allocating OSAM Data Sets | 324 |
| Writing a Load Program | 325 |
| The Load Process | 325 |
| Status Codes for Load Programs | 326 |
| Using SSAs in a Load Program | 327 |
| Loading a Sequence of Segments with the D Command Code | 327 |
| Two Types of Initial Load Program | 327 |
| Basic Initial Load Program | 329 |
| Restartable Initial Load Program | 332 |
| JCL for the Initial Load Program | 336 |
| Loading a HISAM Database | 336 |
| Loading a SHISAM Database | 337 |
| Loading a GSAM Database | 337 |
| Loading an HDAM or a PHDAM Database | 337 |
| Loading a HIDAM or a PHIDAM Database | 337 |
| Loading a Database with Logical Relationships or Secondary Indexes | 337 |
| Loading Fast Path Databases | 337 |
| Loading an MSDB | 337 |
| Loading a DEDB | 338 |
| Loading Sequential Dependent Segments | 339 |
| Chapter 14. Monitoring Databases | 341 |
| IMS Monitor | 341 |
| Monitoring Fast Path Systems | 343 |
| Fast Path Log Analysis Utility | 343 |
| Fast Path Log Reduction | 344 |
| Fast Path Transaction Timings | 344 |
| Monitored Events for Fast Path | 345 |
| Selecting Transactions | 345 |
| Interpreting Fast Path Analysis Reports | 345 |
| Chapter 15. Tuning Databases | 347 |
| Reorganizing the Database | 347 |
| When You Should Reorganize | 348 |
| Reorganizing Databases Offline | 348 |
| Protecting Your Database During an Offline Reorganization | 348 |
| Offline Reorganization Utilities | 349 |
| Partial Offline Reorganization | 349 |
| Offline Reorganization Using UCF | 350 |
| Offline Reorganization Without UCF | 350 |
| HISAM Reorganization Unload Utility (DFSURUL0) | 353 |
| HISAM Reorganization Reload Utility (DFSURRL0) | 354 |
| HD Reorganization Unload Utility (DFSURGU0) | 354 |
| HD Reorganization Reload Utility (DFSURGL0) | 355 |
| Database Prereorganization Utility (DFSURPR0) | 356 |
| Database Scan Utility (DFSURGS0) | 356 |
| Database Prefix Resolution Utility (DFSURG10) | 357 |
| Database Prefix Update Utility (DFSURGP0) | 358 |
| Using HISAM Unload and Reload Utilities for Secondary Indexing Operations | 359 |
| Utility Control Facility (DFSUCF00) | 361 |

| | |
|---|-----|
| Database Surveyor Utility (DFSPRSUR) | 361 |
| Partial Database Reorganization Utility (DFSPRCT1) | 362 |
| Procedures for Offline Database Reorganizations | 363 |
| Reorganizing a HISAM Database (No Secondary Indexes) | 364 |
| Reorganizing an HD (HDAM or HIDAM) Database (No Logical Relationships or Secondary Indexes) | 364 |
| Reorganizing a Primary or Secondary Index | 364 |
| Reorganizing HALDBs | 364 |
| HALDB Offline Reorganization | 365 |
| Overview of HALDB Offline Reorganization | 365 |
| Options for Offline Reorganization of HALDBs | 366 |
| Unloading HALDB Partitions and Databases for Offline Reorganization | 367 |
| Reallocating HALDB Database Data Sets for Offline Reorganization | 368 |
| Reloading HALDB Partitions and Databases for Offline Reorganization | 369 |
| HALDB Online Reorganization | 370 |
| The Initialization Phase for HALDB Online Reorganization | 371 |
| The Copying Phase for HALDB Online Reorganization | 372 |
| The Termination Phase for HALDB Online Reorganization | 374 |
| Migration and Coexistence Considerations for HALDB Online Reorganization | 375 |
| Fallback Considerations for HALDB Online Reorganization | 376 |
| Restrictions for HALDB Online Reorganization | 376 |
| Data Set Naming Conventions for HALDB Online Reorganization | 378 |
| Output Data Set Requirements for HALDB Online Reorganization | 379 |
| Starting HALDB Online Reorganization | 379 |
| Monitoring HALDB Online Reorganization | 380 |
| Modifying and Tuning HALDB Online Reorganization | 380 |
| How a HALDB Online Reorganization Impacts IMS Logging | 383 |
| Controlling the Overall System Impact of a HALDB Online Reorganization | 383 |
| IMS Restart and XRF Processing for HALDB Online Reorganization | 383 |
| IMS Restart and Fast Database Recovery Processing for HALDB Online Reorganization | 384 |
| IMS Remote Site Recovery Processing for HALDB Online Reorganization | 384 |
| Locking Impacts of HALDB Online Reorganization | 384 |
| Using IMS Utilities with HALDB Online Reorganization | 385 |
| Recovery for HALDB Online Reorganization | 386 |
| Activating Sequential Buffering to Improve the Performance of HALDB Online Reorganization | 388 |
| The HALDB Self-Healing Pointer Process | 388 |
| How the Self-Healing Pointer Process Works | 388 |
| Finding Target Segments | 390 |
| Healing Pointers | 391 |
| Performance of the Self-Healing Process | 392 |
| Changing DL/I Access Methods | 394 |
| Changing the DL/I Access Method From HISAM to HIDAM | 395 |
| Changing the DL/I Access Method From HISAM to HDAM | 395 |
| Changing the DL/I Access Method From HIDAM to HISAM | 397 |
| Changing the DL/I Access Method From HIDAM to HDAM | 397 |
| Changing the DL/I Access Method From HDAM to HISAM | 398 |
| Changing the DL/I Access Method From HDAM to HIDAM | 399 |
| Changing the DL/I Access Method From HDAM to PHDAM and HIDAM to PHIDAM | 401 |
| Secondary Index Considerations | 401 |
| Determining the Database Name | 402 |
| Performing the Migration Unload | 402 |
| Performing the Migration Reload | 402 |

| | | |
|---|---|------------|
| | Changing the DL/I Access Method From PHDAM and PHIDAM to HDAM and HIDAM | 402 |
| I | Changing HALDB Partition Definitions | 404 |
| I | Making Changes to a Single Partition | 404 |
| I | Making Changes that Affect All of the Partitions in a HALDB. | 404 |
| I | Updating Partition Structures When Changing Partition Definitions | 405 |
| I | Changing Partition Boundaries. | 406 |
| | Procedure for Changing to DEDBs | 407 |
| | Changing the Hierarchic Structure | 408 |
| | Changing the Sequence of Segment Types | 408 |
| | Combining Segments | 408 |
| | Procedure for Changing the Hierarchic Structure | 409 |
| | Changing Direct-Access Storage Devices. | 409 |
| | Tuning OSAM Sequential Buffering | 410 |
| | Well-Organized Database | 410 |
| | Badly-Organized Database | 410 |
| | Ensuring a Well-Organized Database | 410 |
| | Adjusting HDAM and PHDAM Options | 411 |
| | Adjusting Buffers | 411 |
| | VSAM Buffers | 411 |
| | Monitoring VSAM Buffers | 412 |
| | When to Adjust VSAM Buffers | 412 |
| | VSAM Buffer Adjustment Options. | 412 |
| | OSAM Buffers. | 413 |
| | Procedure for Adjusting VSAM and OSAM Database Buffers | 413 |
| | OSAM Sequential Buffering | 414 |
| | Procedure for Adjusting Sequential Buffers | 414 |
| | Adjusting VSAM Options | 414 |
| | Procedure for Adjusting VSAM Options Specified in the OPTIONS Control Statement | 415 |
| | Procedures for Adjusting VSAM Options Specified in the Access Method Service DEFINE CLUSTER Command | 415 |
| | Changing the FREESPACE Parameter | 415 |
| | Changing the SPEED / RECOVERY Parameter | 416 |
| | Changing the REPLICATE / NOREPLICATE Parameter | 416 |
| | Adjusting OSAM Options. | 416 |
| | Changing the Amount of Space Allocated. | 416 |
| | Changing Operating System Access Methods | 417 |
| | Changing the Number of Data Set Groups | 418 |
| | Tuning Fast Path Systems | 421 |
| | Transaction Volume to a Particular Fast Path Application Program | 422 |
| | DEDB Structure Considerations | 422 |
| | Usage of Buffers from a Buffer Pool. | 423 |
| | Contention for DEDB Control Interval (CI) Resources | 424 |
| | Exhaustion of DEDB DASD Space | 425 |
| | Utilization of Available Real Storage. | 425 |
| | Synchronization Point Processing and Physical Logging | 425 |
| | Contention for Output Threads. | 425 |
| | Overhead Resulting from Reprocessing | 426 |
| | Dispatching Priority of Processor-Dominant and I/O-Dominant Tasks. | 426 |
| | DASD Contention Due to I/O on DEDBs | 426 |
| | Resource Locking Considerations with Block Level Sharing | 426 |
| | Resource Name Hash Routine | 427 |
| | Chapter 16. Modifying Databases | 429 |
| | Adding Segment Types | 430 |

| | |
|---|-----|
| Unloading and Reloading Using the Reorganization Utilities | 430 |
| Without Unloading or Reloading | 431 |
| Using Your Own Unload and Reload Program | 431 |
| Deleting Segment Types | 431 |
| Moving Segment Types | 432 |
| Changing Segment Size | 432 |
| Changing Data in a Segment (Except for Data at the End of a Segment) | 433 |
| Changing the Position of Data in a Segment | 433 |
| Adding Logical Relationships | 433 |
| Example 1. DBX Exists, DBY Is to Be Added | 434 |
| Example 1 Procedure | 434 |
| Example 2. DBX and DBY Exist, DBZ Is to Be Added | 435 |
| Example 2 Procedure | 435 |
| Example 3. DBX and DBY Exist, DBZ Is to Be Added | 436 |
| Example 4. DBX and DBY Exist, DBZ Is to Be Added | 436 |
| Example 5. DBX Exists, DBY Is to Be Added | 437 |
| Procedure | 437 |
| Example 6. DBX and DBY Exist, DBZ Is to Be Added | 438 |
| Procedure When Reorganizing DBY (Segment B Contains a Symbolic Pointer) | 438 |
| Procedure When Reorganizing DBY and Scanning DBX (Segment B Contains a Direct Pointer) | 439 |
| Procedure When Reorganizing DBX and DBY | 439 |
| Example 7. DBX and DBY Exist, DBZ Is to Be Added | 440 |
| Procedure Using Scan | 440 |
| Procedure When Reorganizing DBX and DBY | 441 |
| Example 8. DBX and DBY Exist, DBZ Is to Be Added | 442 |
| Example 9. DBY Exists, DBZ Is to Be Added | 442 |
| Procedure | 442 |
| Example 10. DBY Exists, DBZ Is to Be Added | 443 |
| Example 11. DBX and DBY Exist, DBZ Is to Be Added. | 443 |
| Example 12. DBX and DBY Exist, DBZ Is to Be Added. | 444 |
| Example 13. DBX and DBY Exist, Segment Y and DBZ Are to Be Added | 444 |
| Procedure | 444 |
| Steps in Reorganizing a Database to Add a Logical Relationship | 445 |
| Some Restrictions on Modifying Existing Logical Relationships | 449 |
| Example 1: Changing from Bidirectional Virtual to Bidirectional Physical Pairing | 449 |
| Example 2: Changing the Location of the Real Logical Child in a Bidirectional Logical Relationship | 450 |
| Summary on Use of Utilities When Adding Logical Relationships | 450 |
| Adding a Secondary Index | 451 |
| Adding or Converting to Variable-Length Segments | 451 |
| Method 1. Converting Segments or a Database | 451 |
| Method 2. Converting Segments or a Database | 452 |
| Converting to the Segment Edit/Compression Exit Routine | 452 |
| Converting Databases for Data Capture Exit Routines and Asynchronous Data Capture | 453 |
| Converting a Logical Parent Concatenated Key from Virtual to Physical or Physical to Virtual | 454 |
| Using the Online Change Function | 454 |
| Online Change, DEDBs, and Availability of IFP and MPP Regions | 455 |
| Online Change and DEDB Randomizer and Exit Routines | 457 |
| Introducing a New DEDB Randomizing Routine Using Online Change | 457 |
| Changing DEDB Randomizing Routines Using Online Change | 458 |
| Deleting a DEDB Randomizing Routine Using Online Change | 458 |

| | |
|---|-----|
| Online Change and DEDB Segment Compression Routines | 458 |
| Online Change and DEDB Data Capture Exit Routines. | 459 |
| Changing DEDB Root Addressable Space with Two Stage Randomizer Using Online Change | 460 |
| Changing the DEDB AREA UOW Structural Definition Using Online Change | 460 |
| Making Online Changes at the DEDB and Area Level | 461 |
| Adding or Deleting DEDBs Using Online Change | 461 |
| Adding or Deleting DEDB Segments Using Online Change | 462 |
| Adding or Deleting DEDB Areas Using Online Change | 463 |
| Changing DEDB Root Addressable Space Allocation Using Online Change | 463 |
| Online Change and DEDB Dependent and Independent Overflow Space Allocation | 464 |
| Changing DEDB CI Size Using Online Change | 464 |
| Extending DEDB Independent Overflow Online | 464 |

Chapter 5. Analyzing Data Requirements

One of the early steps of database design is developing a conceptual data structure that satisfies your end user's processing requirements. So, before you can develop a conceptual data structure, familiarize yourself with your end user's processing and data requirements.

Developing a data structure is a process of combining the data requirements of each of the tasks to be performed, into one or more data structures that satisfy those requirements. The method explained here describes how to use the local views developed for each business process to develop a data structure.

A business process, in an application, is one of the tasks your end user needs done. For example, in an education application, printing a class roster is a business process.

A local view describes a conceptual data structure and the relationships between the pieces of data in the structure for one business process.

To understand the method explained in this chapter, you need to be familiar with the terminology and examples explained in the introductory chapter on application design in *IMS Version 9: Application Programming: Design Guide*. The chapter of the design guide explains how to develop local views for the business processes in an application.

Included in this chapter are the following topics:

- "Local View," which introduces you to the local view examples and explains the information that makes up a local view.
- "Designing a Conceptual Data Structure" on page 53, which explains how you can develop a conceptual data structure based on the local views for the business processes in an application.
- "Implementing the Structure with DL/I" on page 55, which explains how you implement the structure you have developed with DL/I. The considerations explained are: assigning data elements to segments and resolving data conflicts with DL/I.

Local View

Designing a structure that satisfies the data requirements of the business processes in an application requires an understanding of the requirements for each of those business processes. A local view of the business process describes these requirements because the local view provides:

- A list of all the data elements the process requires and their controlling keys
- The conceptual data structure developed for each process, showing how the data elements are grouped into data aggregates
- The mappings between the data aggregates in each process

This chapter uses a company that provides technical education to its customers as an example. The education company has one headquarters, called HQ, and several local education centers, called Ed Centers. HQ develops the courses offered at each of the Ed Centers. Each Ed Center is responsible for scheduling classes it will offer and for enrolling students for those classes.

Local View

A class is a single offering of a course on a specific date at an Ed Center. There might be several offerings of one course at different Ed Centers, and each of these offerings is a separate class.

The local views used in this chapter are for the following business processes in an education application:

- Current Roster
- Schedule of Classes
- Instructor Skills Report
- Instructor Schedules

The information in the subtopics of this topic summarizes the local views developed in the introductory chapter on application design in *IMS Version 9: Application Programming: Design Guide*.

Notes for local views:

- The asterisks (*) in the data structures for each of the local views indicate the data elements that identify the data aggregate. This is the data aggregate's key; some data aggregates require more than one data element to uniquely identify them.
- The mappings between the data aggregates in each process are given in mapping notation. A one-to-many mapping means for each *A* aggregate there are one or more *B* aggregates; shown like this: \longleftrightarrow

A many-to-many relationship means that for each *A* aggregate there are many *B* aggregates, and for each *B* aggregate, there are many *A* aggregates; shown as follows: \longleftrightarrow

Local View 1. Current Roster

This topic describes the elements, the data structure, the data aggregates, and the mapping of the relationships between the data aggregates used to satisfy the data requirements of the Current Roster business process.

List of Current Roster Data Elements

The following is a list of the data elements and their descriptions for our technical education provider example.

| Data Element | Description |
|--------------|-----------------------------|
| CRSNAME | Course name |
| CRSCODE | Course code |
| LENGTH | Length of class |
| EDCNTR | Ed Center offering class |
| DATE | Date class is offered |
| CUST | Customer that sent student |
| LOCTN | Location of customer |
| STUSEQ# | Student's sequence number |
| STUNAME | Student's name |
| STATUS | Student's enrollment status |
| ABSENCE | Student's absences |

| | |
|---------------|---------------------------|
| GRADE | Student's grade for class |
| INSTRS | Instructors for class |

Figure 17 shows the conceptual data structure for the current roster.

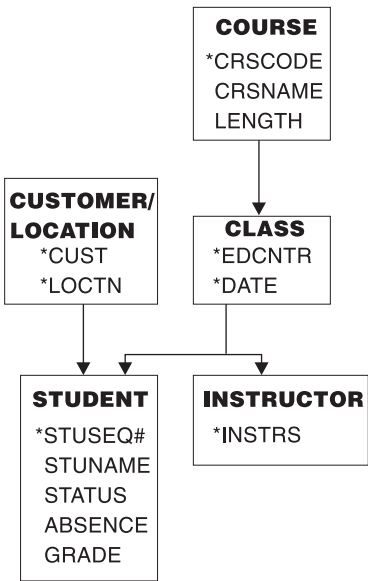


Figure 17. Current Roster Conceptual Data Structure

Current Roster Mappings

The mappings for the current roster are:

- Course ↔ Class
- Class ↔ Student
- Class ↔ Instructor
- Customer/location ↔ Student

Local View 2. Schedule of Classes

| This topic describes the elements, the data structure, the data aggregates, and the
| mapping of the relationships between the data aggregates used to satisfy the data
| requirements of the Schedule of Classes business process.

List of Schedule of Classes Data Elements

The following is a list of the schedule of classes and their descriptions for our example.

| Data Element | Description |
|----------------|---|
| CRSCODE | Course code |
| CRSNAME | Course name |
| LENGTH | Length of course |
| PRICE | Price of course |
| EDCNTR | Ed Center where class is offered |
| DATE | Dates when class is offered at a particular Ed Center |

Local View

Figure 18 shows the conceptual data structure for the class schedule.




Figure 18. Schedule of Classes Conceptual Data Structure

Schedule of Classes Mappings

The only mapping for this local view is:

Course \longleftrightarrow Class

Local View 3. Instructor Skills Report

This topic describes the elements, the data structure, the data aggregates, and the mapping of the relationships between the data aggregates used to satisfy the data requirements of the Instructor Skills Report business process.

List of Instructor Skills Report Data Elements

The following is a list of the instructor skills report data elements and their descriptions for our technical education provider example.

| Data Element | Description |
|--------------|-------------|
| INSTR | Instructor |
| CRSCODE | Course code |
| CRSNAME | Course name |

Figure 19 shows the conceptual data structure for the instructor skills report.




Figure 19. Instructor Skills Report Conceptual Data Structure

Instructor Skills Report Mappings

The only mapping for this local view is:

Instructor \longleftrightarrow Course

Local View 4. Instructor Schedules

| This topic describes the elements, the data structure, the data aggregates, and the
| mapping of the relationships between the data aggregates used to satisfy the data
| requirements of the Instructor Schedules business process.

List of Instructor Schedules Data Elements

The following is a list of the instructor schedules data elements and their descriptions for our example.

| Data Element | Description |
|--------------|----------------------------|
| INSTR | Instructor |
| CRSNAME | Course name |
| CRSCODE | Course code |
| EDCNTR | Ed Center |
| DATE | Date when class is offered |

Figure 20 shows the conceptual data structure for the instructor schedules.





Figure 20. Instructor Schedules Conceptual Data Structure

Instructor Schedules Mappings

The mappings for this local view are:



Designing a Conceptual Data Structure

Analyzing the mappings from all the local views is one of the first steps in designing a conceptual data structure. Two kinds of mappings affect the segments: one-to-many and many-to-many.

A one-to-many mapping means that for each segment *A* there are one or more segment *B*s; shown like this: *A* \longleftrightarrow *B*. For example, in the Current Roster (Figure 17 on page 51), there is a one-to-many relationship between course and class. For each course, there can be several classes scheduled, but a class is

Designing a Conceptual Data Structure

associated with only one course. A one-to-many relationship can be represented as a dependent relationship: In the course/class example, the classes are dependent on a particular course.

A many-to-many mapping means that for each segment *A* there are many segment *B*s, and for each segment *B* there are many segment *A*s. This is shown like this: $A \longleftrightarrow B$. A many-to-many relationship is not a dependent relationship, since it usually occurs between data aggregates in two separate data structures and indicates a conflict in the way two business processes need to process that data.

When you implement a data structure with DL/I, there are three strategies you can apply to solve data conflicts:

- Defining logical relationships

- Establishing secondary indexes

- Storing the data in two places (also known as carrying duplicate data).

Related Reading: “Resolving Data Conflicts” on page 56 explains the kinds of data conflicts that secondary indexes and logical relationships can resolve.

The first step in designing a conceptual data structure is to combine the mappings of all the local views. To do this, go through the mappings for each local view and make a consolidated list of mappings (see Table 5). As you review the mappings:

- Do not record duplicate mappings. At this stage you need to cover each variation, not each occurrence.
- If two data aggregates in different local views have opposite mappings, use the more complex mapping. This will include both mappings when they are combined. For example, if local view #1 has the mapping $A \longleftrightarrow B$, and local view #2 has the mapping $A \longleftarrow B$, use a mapping that includes both these mappings. In this case, this is $A \longleftrightarrow B$.

Table 5. Combined Mappings for Local Views

| Mapping | Local View |
|---|------------|
| Course \longleftrightarrow Class | 1, 2, 4 |
| Class \longleftrightarrow Student | 1 |
| Class \longleftrightarrow Instructor | 1 |
| Customer/location \longleftrightarrow Student | 1 |
| Instructor \longleftrightarrow Course | 3, 4 |

Using the combined mappings, you can construct the data structures shown in Figure 21.




Figure 21. Education Data Structures

Two conflicts exist in these data structures. First, STUDENT is dependent on both CUST and CLASS. Second, there is an opposite mapping between COURSE and INSTR, and INSTR and COURSE. If you implemented these structures with DL/I, you could use logical relationships to resolve the conflicts. “Analyzing Requirements for Logical Relationships” on page 56 explains how.

Implementing the Structure with DL/I

When you implement a data structure with DL/I, you implement it as a hierarchy. A hierarchy is made up of segments. In a hierarchy, a one-to-many relationship is called a parent/child relationship. In a hierarchy, each segment can have one or more children, but it can have only one parent.

When you use DL/I, consider how each of the data elements in the structure you have developed should be grouped into segments. Also, consider how DL/I can solve any existing data conflicts in the structure. The topics “Assigning Data Elements to Segments” and “Resolving Data Conflicts” on page 56 in this chapter explain how you assign data elements to segments, and how DL/I can resolve data conflicts.

Assigning Data Elements to Segments

Once you determine how data elements are related in a hierarchy, associate each of the data elements with a segment. To do this, construct a list of all the keys and their associated data elements. If a key and its associated data element appear in several local views, only record the association once.

List the data elements next to their keys, as shown in Table 6. The key and its associated data elements become the segment content.

Table 6. Keys and Associated Data Elements

| Data Aggregate | Key | Data Elements |
|-------------------|--------------|---------------------------------|
| COURSE | CRSCODE | CRSNAME, LENGTH, PRICE |
| CUSTOMER/LOCATION | CUST, LOCTN | |
| CLASS | EDCNTR, DATE | |
| STUDENT | STUSEQ# | STUNAME, ABSENCE, STATUS, GRADE |
| INSTRUCTOR | INSTR | |

If a data element is associated with different keys in different local views, then you must decide which segment will contain the data element. The other thing you can do is to store duplicate data. To avoid doing this, store the data element with the key that is highest in the hierarchy. For example, if the keys ALPHA and BETA were both associated with the data element XYZ (one in local view 1 and one in local view 2), and ALPHA were higher in the hierarchy, store XYZ with ALPHA to avoid having to repeat it.

Resolving Data Conflicts

The data structure you design can fall short of the application's processing requirements. For example, one business process might need to retrieve a particular segment by a field other than the one you have chosen as the key field. Another business process might need to associate segments from two or more different data structures. Once you have identified these kinds of conflicts in a data structure and are using DL/I, you can look at two DL/I options that can help you resolve the conflicts: secondary indexing and logical relationships.

Analyzing Requirements for Secondary Indexes

Secondary indexing allows a segment to be identified by a field other than its key field.

Suppose that you are part of our technical education company and need to determine (from a terminal) whether a particular student is enrolled in a class. If you are unsure about the student's enrollment status, you probably do not know the student's sequence number. The key of the STUDENT segment, however, is STUSEQ#. Let's say you issue a request for a STUDENT segment, and identify the segment you need by the student's name (STUNAME). Instead of the student's sequence number (STUSEQ#), IMS searches through all STUDENT segments to find that one. Assuming the STUDENT segments are stored in order of student sequence numbers, IMS has no way of knowing where the STUDENT segment is just by having the STUNAME.

Using a secondary index in this example is like making STUNAME the key field of the STUDENT segment for this business process. Other business processes can still process this segment with STUSEQ# as the key.

To do this, you can index the STUDENT segment on STUNAME in the secondary index. You can index any field in a segment. When you index a field, indicating to IMS that you are using a secondary index for that segment, IMS processes the segment as though the indexed field were the key.

Analyzing Requirements for Logical Relationships

When a business process needs to associate segments from different hierarchies, logical relationships can make that possible.

Defining logical relationships lets you create a hierarchic structure that does not exist in storage but can be processed as though it does. You can relate segments in separate hierarchies. The data structure created from these logical relationships is called a logical structure. To relate segments in separate hierarchies, store the segment in the path by which it is accessed most frequently. Store a pointer to the segment in the path where it is accessed less frequently.

In the hierarchy shown in Figure 21 on page 55, two possible parents exist for the STUDENT segment. If the CUST segment is part of an existing database, you can

define a logical relationship between the CUST segment and the STUDENT segment. You would then have the hierarchies shown in Figure 22. The CUST/STUDENT hierarchy would be a logical structure.

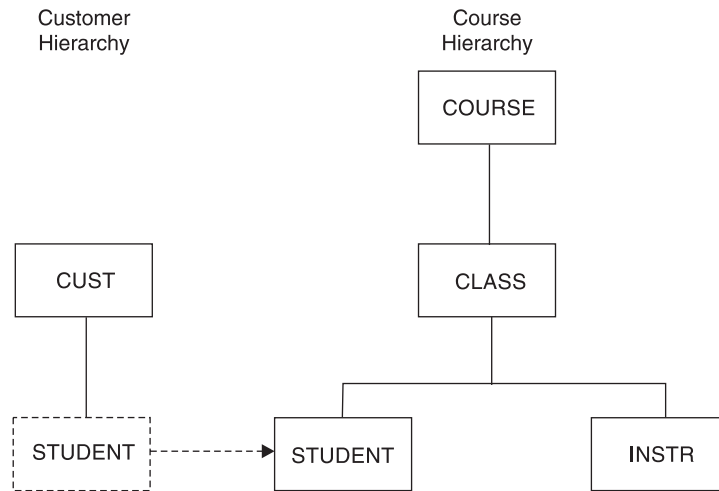


Figure 22. Education Hierarchies

This kind of logical relationship is called unidirectional, because the relationship is “one way.”

The other conflict you can see in Figure 21 on page 55, is the one between COURSE and INSTR. For one course there are several classes, and for one class there are several instructors (COURSE \longleftrightarrow CLASS \longleftrightarrow INSTR), but each instructor can teach several courses (INSTR \longleftrightarrow COURSE). You can resolve this conflict by using a bidirectional logical relationship. You can store the INSTR segment in a separate hierarchy, and store a pointer to it in the INSTR segment in the course hierarchy. You can also store the COURSE segment in the course hierarchy, and store a pointer to it in the COURSE segment in the INSTR hierarchy. This bidirectional logical relationship would give you the two hierarchies shown in Figure 23, eliminating the need to carry duplicate data.

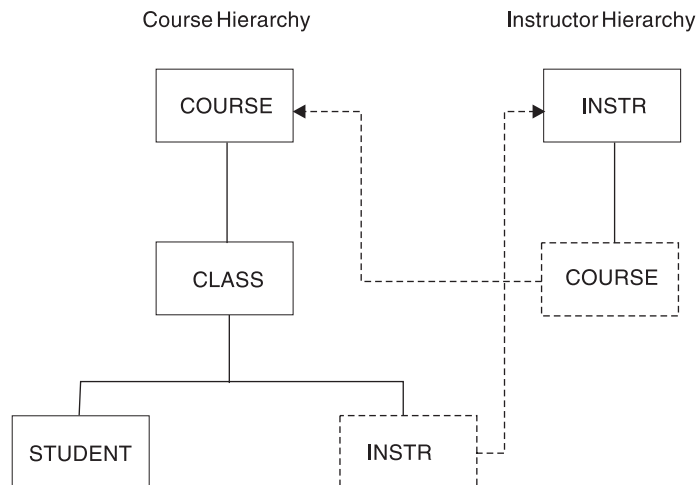


Figure 23. Bidirectional Logical Relationships

Chapter 6. Choosing Full-Function Database Types

IMS databases are hierarchic databases that are accessed through DL/I calls. IMS makes it possible for application programs to retrieve, replace, delete, and add segments to IMS databases.

IMS allows you to define twelve database types. Each type has different organization processing characteristics. Except for DEDB and MSDB, all the database types are discussed in this chapter.

In this Chapter:

- “Sequential Storage Method” on page 60
- “Direct Storage Method” on page 60
- “Databases Supported with DBCTL” on page 60
- “Databases Supported with DCCTL” on page 61
- “Performance Considerations Overview” on page 61
- “HSAM Databases” on page 64
- “HISAM Databases” on page 69
- “SHSAM, SHISAM and GSAM Databases” on page 78
- “HDAM, PHDAM, HIDAM, and PHIDAM Databases” on page 82
- “Managing I/O Errors” on page 111

Related Reading: For information on DEDBs and MSDBs see, “Data Entry Databases” on page 113 and “Main Storage Databases (MSDBs)” on page 132.

Understanding how the database types differ enables you to pick the type that best suits your application’s processing requirements.

Each database type has its own access method. The following figure lists each type and the access method it uses:

| Type of Database | Access Method |
|------------------|--|
| HSAM | Hierarchical Sequential Access Method |
| HISAM | Hierarchical Indexed Sequential Access Method |
| SHSAM | Simple Hierarchical Sequential Access Method |
| SHISAM | Simple Hierarchical Indexed Sequential Access Method |
| GSAM | Generalized Sequential Access Method |
| | <u>Restriction:</u> GSAM does not apply to CICS applications. |
| HDAM | Hierarchical Direct Access Method |
| PHDAM | Partitioned Hierarchical Direct Access Method |
| HIDAM | Hierarchical Indexed Direct Access Method |
| PHIDAM | Partitioned Hierarchical Indexed Direct Access Method |
| PSINDEX | Partitioned Secondary Index Database |
| DEDB | Data Entry Database (Hierarchical Direct Access) |

Sequential Storage Method

MSDB

Main Storage Database (Hierarchical Direct Access)

Based on the access method used, the various databases can be classified into two groups: sequential storage and direct storage.

Sequential Storage Method

HSAM, HISAM, SHSAM, and SHISAM use the sequential method of accessing data. With this method, the hierarchic sequence of segments in the database is maintained by putting segments in storage locations that are physically adjacent to each other. GSAM databases also use the sequential method of accessing data, but no concept of hierarchy, database record, or segment exists in GSAM databases.

Direct Storage Method

HDAM, PHDAM, HIDAM, DEDB, MSDB, and PHIDAM databases use the direct method of accessing data. With this method, the hierarchic sequence of segments is maintained by putting direct-address pointers in each segment's prefix.

For quick reference, see Table 7 on page 63 for a summary of HSAM, HISAM, HDAM, PHDAM, HIDAM, PHIDAM, DEDB, and MSDB database characteristics.

Databases Supported with DBCTL

Database Control (DBCTL) configuration of IMS supports all IMS full-function databases:

HSAM
HISAM
SHSAM
SHISAM
HDAM
PHDAM
HIDAM
PHIDAM
PSINDEX

Databases can be accessed through DBCTL from IMS BMP regions, as well as from independent transaction-management subsystems. Only batch-oriented BMP programs are supported because DBCTL provides no message or transaction support.

CICS online programs can access the same IMS database concurrently; however, an IMS batch program must have exclusive access to the database (if you are not participating in IMS data sharing).

If you have batch jobs that currently access IMS databases through IMS data sharing, you can convert them to run as BMPs directly accessing databases through DBCTL, thereby improving performance. You can additionally convert current batch programs to BMPs to access DEDBs.

Related Reading: For more information on converting a batch job to a BMP, see *IMS Version 9: Application Programming: Design Guide* and *IMS Version 9: Administration Guide: System*.

Databases Supported with DCCTL

The DCCTL configuration of IMS supports the following database and dependent region combinations:

- GSAM databases for BMP regions
- DB2 UDB for z/OS databases for BMP, MPP, and IFP regions through the External Subsystem attachment facility (ESAF)
- DB2 UDB for z/OS databases for JMP and JBP regions through the DB2 Recoverable Resource Manager Services attachment facility (RRSAF)

Restriction: DCCTL does not support full-function or Fast Path databases.

Related Reading:

- For more information on ESAF, see *IMS Version 9: Customization Guide*
- For more information on RRSAF, see *DB2 Universal Database for z/OS Administration Guide*

Performance Considerations Overview

All databases are not created equal. You will want to make an informed decision regarding the type of database organizations which will best serve your purposes. The following lists briefly summarize the performance characteristics of the various database types, highlighting efficiencies and deficiencies of hierarchic sequential, hierarchic direct and general sequential databases.

Related Reading: For information on DEDBs and MSDBs, see “Data Entry Databases” on page 113 and “Main Storage Databases (MSDBs)” on page 132.

General Sequential (GSAM)

- Supported by DCCTL
- No hierarchy, database records, segments, or keys
- No DLET or REPL
- ISRT adds records at end of data set
- GN and GU processed in batch or BMP applications only
- Allows IMS symbolic checkpoint calls and restart from checkpoint (except VSAM-loaded databases)
- Good for converting data to IMS and for passing data
- Not accessible from an MPP or JMP region
- Space efficient
- Not time efficient

VSAM

- Fixed- or variable-length records are usable
- VSAM ESDS DASD stored
- IMS symbolic checkpoint call allowed
- Restart from checkpoint not allowed

BSAM/QSAM

Performance Considerations Overview

- Fixed-, variable-, or undefined-length records are usable
- BSAM/QSAM DS tape or DASD stored
- Allows IMS symbolic checkpoint calls and restart from checkpoint

Hierarchic Sequential

Segments are linked by physical contiguity

HSAM

- Supported by DBCTL
- Physical sequential access to roots and dependents stored on tape or DASD
- ISRT allowed only when database is loaded
- GU, GN, and GNP allowed
- Database update done by merging databases and writing new database
- QSAM and BSAM accessible
- Space efficient but not time efficient
- Sequential access

HISAM

- Supported by DBCTL
- Hierarchic indexed access to roots
- Sequential access to dependents
- Stored on DASD
- VSAM accessible
- All DL/I calls allowed
- Index is on root segment sequence field
- Good for databases not updated often
- Not space efficient with many updates
- Time efficient with SSA-qualified calls

SHSAM

- Supported by DBCTL
- Simple hierarchic sequential access method to root segments only
- ISRT allowed only when database is loaded
- GU, GN, and GNP allowed
- Database update done by reloaded database
- QSAM and BSAM accessible
- Allows IMS symbolic checkpoint calls and restart from checkpoint (except VSAM-loaded databases)
- Good for converting data to IMS and for passing data
- Not accessible from an MPP or JMP region
- Space efficient
- Not time efficient

SHISAM

- Supported by DBCTL
- Simple hierarchic indexed access to roots only
- Sequential access to dependents

- Stored on DASD
- VSAM accessible
- All DL/I calls allowed
- Good for converting data to IMS and for passing data
- Not space efficient
- Time efficient

Hierarchic Direct

Segments are linked by pointers

HDAM and PHDAM

- Supported by DBCTL
- Hashing access to roots
- Sequential access by secondary index to segments
- All DL/I calls allowed
- Stored on DASD in VSAM ESDS or OSAM data set
- Good for direct access to records
- Hierarchic pointers allowed
 - Hierarchic sequential access to dependent segments
 - Better performance than child and twin pointers
 - Less space required than child and twin pointers
- Child and twin pointers allowed
 - Direct access to pointers
 - More space required by additional index VSAM ESDS database

HIDAM and PHIDAM

- Supported by DBCTL
- Indexed access to roots
- Pointer access to dependent segments
- All DL/I calls allowed
- Stored on DASD in VSAM ESDS or OSAM data set
- Good for random and sequential access to records
- Good for random access to segment paths
- Hierarchic pointers allowed
 - Hierarchic sequential access to dependent segments
 - Better performance than child and twin pointers
 - Less space required than child and twin pointers
- Child and twin pointers allowed
 - Direct access to pointers
 - More space required by additional index VSAM ESDS database

Table 7 gives a summary of database characteristics, functions, and options for the different database types.

Table 7. Summary of Database Characteristics and Options for Database Types

| Characteristic | HSAM | HISAM | HDAM | PHDAM | HIDAM | PHIDAM | DEDB | MSDB |
|-------------------------|------|-------|------|-------|-------|--------|------|------|
| Hierarchical Structures | Y | Y | Y | Y | Y | Y | Y | N |

Performance Considerations Overview

Table 7. Summary of Database Characteristics and Options for Database Types (continued)

| Characteristic | HSAM | HISAM | HDAM | PHDAM | HIDAM | PHIDAM | DEDB | MSDB |
|---|------|-------|------|-------|-------|--------|------|------|
| Direct Access Storage | Y | Y | Y | Y | Y | Y | Y | N |
| Multiple Data Set Groups | N | N | Y | Y | Y | Y | N | N |
| Logical Relationships | N | Y | Y | Y | Y | Y | N | N |
| Variable-Length Segments | N | Y | Y | Y | Y | Y | Y | N |
| Segment Edit/Compression | N | Y | Y | Y | Y | Y | Y | N |
| Data Capture Exit Routines | N | Y | Y | Y | Y | Y | Y | N |
| Field-Level Sensitivity | Y | Y | Y | Y | Y | Y | N | N |
| Primary Index | N | Y | N | N | Y | Y | N | N |
| Secondary Index | N | Y | Y | Y | Y | Y | N | N |
| Logging, Recovery, Offline Reorganization | N | Y | Y | Y | Y | Y | Y | Y |
| VSAM | N | Y | Y | Y | Y | Y | Y | N/A |
| OSAM | N | N | Y | Y | Y | Y | N | N/A |
| QSAM/BSAM | Y | N | N | N | N | N | N | N/A |
| Boolean Operators | Y | Y | Y | Y | Y | Y | Y | N |
| Command Codes | Y | Y | Y | Y | Y | Y | Y | N |
| Subset Pointers | N | N | N | N | N | N | Y | N |
| Uses Main Storage | N | N | N | N | N | N | N | Y |
| High Parallelism (field call) | N | N | N | N | N | N | N | Y |
| Compaction | Y | Y | Y | Y | Y | Y | Y | N |
| DBRC Support | Y | Y | Y | Y | Y | Y | Y | N/A |
| Partitioning Support | N | N | N | Y | N | Y | Y | N |
| Data Sharing | Y | Y | Y | Y | Y | Y | Y | N |
| Partition Sharing | N | N | N | Y | N | Y | Y | N |
| Block Level Sharing | Y | Y | Y | Y | Y | Y | Y | N |
| Area Sharing | N/A | N/A | N/A | N/A | N/A | N/A | Y | N/A |
| Record Deactivation | N | N | N | N | N | N | Y | N/A |
| Database Size | med | med | med | lg | med | lg | lg | sml |
| Online Utilities | N | N | N | N | N | N | Y | N |
| Online Reorganization | N | N | N | Y | N | Y | Y | N |
| Batch | Y | Y | Y | Y | Y | Y | N | N |

HSAM Databases

Hierarchical sequential access method (HSAM) databases use the sequential method of accessing data. All database records and all segments within each database record are physically adjacent in storage. An HSAM database can be stored on tape or on a direct-access storage device. They are processed using

either basic sequential access method (BSAM) or queued sequential access method (QSAM) as the operating system access method. Specify your access method on the PROCOPT= parameter in the PCB. If you specify PROCOPT=GS, QSAM is always used. If you specify PROCOPT=G, BSAM is used.

HSAM data sets are loaded with root segments in ascending key sequence (if keys exist for the root) and dependent segments in hierarchic sequence. You do not need to define a key field in root segments. You must, however, present segments to the load program in the order in which they must be loaded. HSAM data sets use a fixed-length, unblocked record format (RECFM=F), which means that the logical record length is the same as the physical block size.

HSAM databases can only be updated by rewriting them. Delete (DLET) and replace (REPL) calls are not allowed, and insert (ISRT) calls are only allowed when the database is being loaded. Although the field-level sensitivity option can be used with HSAM databases the following options *cannot* be used with HSAM databases:

- Multiple data set groups
- Logical relationships
- Secondary indexing
- Variable-length segments
- Segment edit/compression exit routine
- Data Capture exit routines
- Asynchronous data capture
- Logging, recovery, or reorganization

Multiple positioning and multiple PCBs cannot be used in HSAM databases.

When to Use HSAM

Although the uses of HSAM are limited because of its processing characteristics, it is used for applications requiring sequential processing only. Typically, HSAM is used for low-use files. These are files containing, for example, audit trails, statistical reports or files containing historical or archive data that has been purged from the main database.

How an HSAM Record Is Stored

Segments in an HSAM database are loaded in the order in which you present them to the load program. You should present all segments within a database record in hierarchic sequence. If a sequence field has been defined for root segments, you should present database records to the load program in ascending root key sequence.

Figure 24 on page 66 shows an example HSAM database.

HSAM Databases

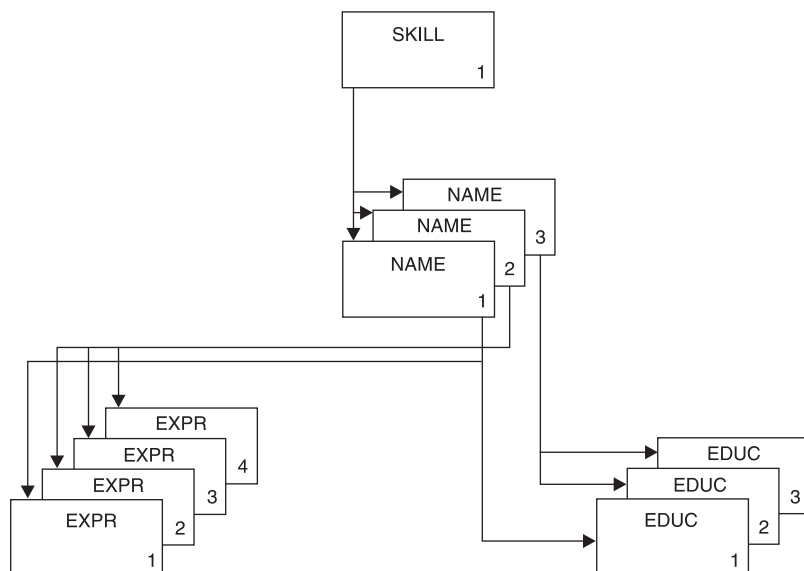


Figure 24. Example HSAM Database

Figure 25 shows how the example HSAM database, shown in Figure 24, would be stored in blocks.

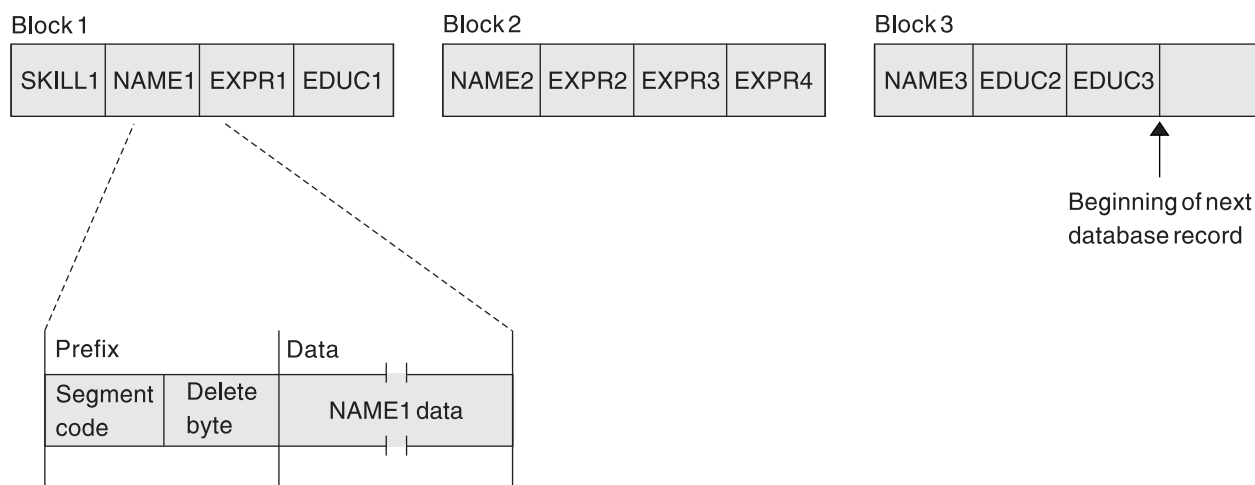


Figure 25. Example HSAM Database Stored in Blocks

In the data set, a database record is stored in one or more consecutive blocks. You define what the block size will be. Each block is filled with segments of the database record until there is not enough space left in the block to store the next segment. When this happens, the remaining space in the block is padded with zeros and the next segment is stored in the next consecutive block. When the last segment of a database record has been stored in a block, any unused space, if sufficient, is filled with segments from the next database record.

In storage, an HSAM segment consists of a 2-byte prefix followed by user data. The first byte of the prefix is the segment code, which identifies the segment *type* to IMS. This number can be from 1 to 255. The segment code is assigned to the segment by IMS in ascending sequence, starting with the root segment and

continuing through all dependents in hierarchic sequence. The second byte of the prefix is the delete byte. Because DLET calls cannot be used against an HSAM database, the second byte is not used.

DL/I Calls against an HSAM Database

Initial entry to an HSAM database is through GU or GN calls. When the first call is issued, the search for the desired segment starts at the beginning of the database and passes sequentially through all segments stored in the database until the desired segment is reached. After the desired segment is reached, its position is used as the starting position for any additional calls that process the database in a forward direction.

After position in an HSAM database has been established, the way in which GU calls are handled depends on whether a sequence field is defined for the root segment and what processing options are in effect. Figure 26 shows a flow chart of the actions taken based on whether a sequence field is defined and what processing options are in effect.

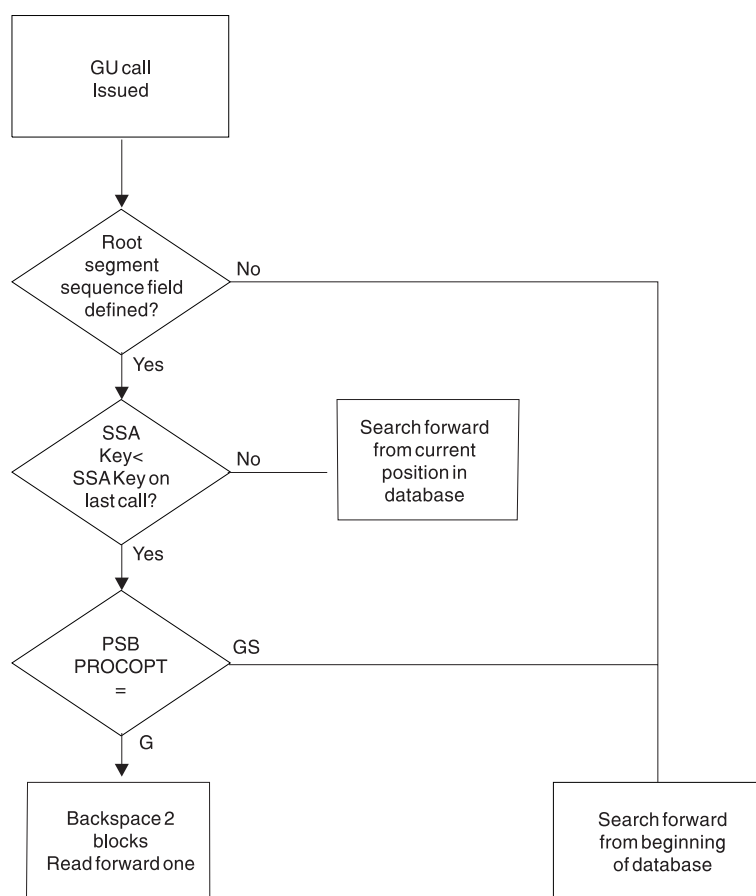


Figure 26. GU Calls against an HSAM Database

When a GU call is issued and the root segment sequence field is not defined, search forward from beginning of database. If the sequence field is defined for the root and the SSA key is less than the SAA[®] key on the last call, search forward from the current position in the database. If the sequence field is defined for the root and the SSA key is greater than the SSA key on the last call, the GU call is

handled based on the PSB PROCOPT. If PROCOPT=GS, search forward from beginning of database. If PROCOPT=G, Backspace two blocks and read forward one block.

No Sequence Field Defined

If no sequence field has been defined, each GU call causes the search for the desired segment to start at the beginning of the database regardless of current position. This allows direct processing of the HSAM database. The processing, however, is restricted to one volume.

Sequence Field Defined

If a sequence field has been defined and the GU call retrieves a segment that is forward in the database, the search starts from the current position and moves forward to the desired segment. If access to the desired segment requires backward movement in the database, the PROCOPT= parameters G or GS (specified during PSBGEN) determine how backward movement is accomplished. If you specify PROCOPT=GS, (that is, the database is read using QSAM), the search for the desired segment starts at the beginning of the database and moves forward. If you specify PROCOPT=G, (that is, the database is read using BSAM), the search moves backward in the database. This is accomplished by backspacing over the block just read and the block previous to it, then reading this previous block forward until the desired segment is found.

Because of the way in which segments are accessed in an HSAM database, it is most practical to access root segments sequentially and dependent segments in hierarchic sequence within a database record. Other methods of access, involving backspacing, rewinding of the tape, or scanning the data set from the beginning, can be time consuming.

As stated previously, DLET and REPL calls cannot be issued against an HSAM database. ISRT calls are allowed only when the database is being loaded. To update an HSAM database, you must write a program that merges the current HSAM database and the update data. The update data can be in one or more files. The output data set created by this process is the new updated HSAM database. Figure 27 illustrates this process.

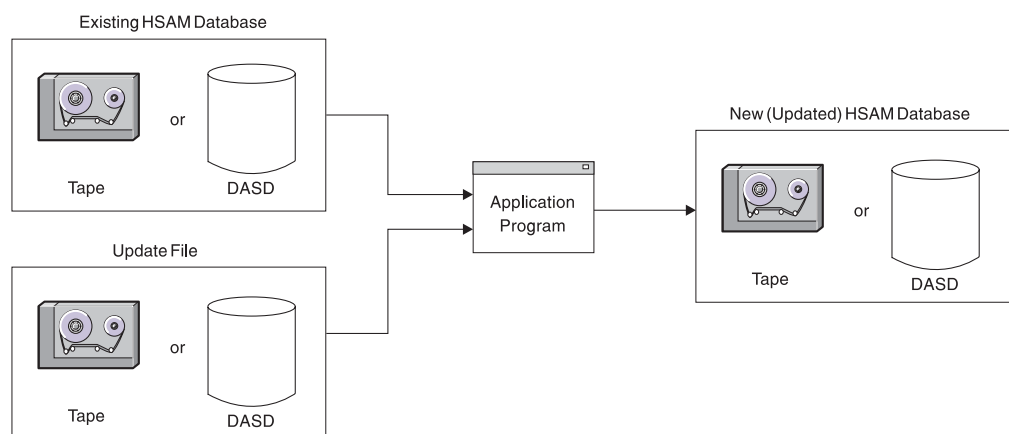


Figure 27. Updating an HSAM Database

HISAM Databases

In a hierarchical indexed sequential access method (HISAM) database, as with an HSAM database, segments in each database record are related through physical adjacency in storage. Unlike HSAM, however, each HISAM database record is indexed, allowing direct access to a database record. In defining a HISAM database, you must define a unique sequence field in each root segment. These sequence fields are then used to construct an index to root segments (and therefore database records) in the database.

HISAM databases are stored on direct-access devices. They can be processed using the virtual storage access method (VSAM) utility. Unlike HSAM, all DL/I calls can be issued against a HISAM database. In addition, the following options are available for HISAM databases:

- Logical relationships
- Secondary indexing
- Variable-length segments
- Segment edit/compression exit routine
- Data Capture exit routines
- Field-level sensitivity
- Logging, recovery, and reorganization

Except for logging and recovery, each of these options is discussed in detail in later parts of this book. For detailed discussions of logging and recovery, see the *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

When to Use HISAM

HISAM is typically used for databases that require direct access to database records and sequential processing of segments in a database record. It is a good candidate for databases with the following characteristics:

- Most database records are about the same size.
- The database does not consist of relatively few root segments and a large number of dependent segments.
- Applications do not depend on a heavy volume of root segments being inserted after the database is initially loaded.
- Deletion of database records is minimal.

More detailed information on the uses of HISAM, requiring a working knowledge of how a HISAM database is organized and processed, is under “Variable-Length Segments” on page 215.

How a HISAM Record is Stored

HISAM database records are stored in two data sets. The first data set, called the *primary data set*, contains an index and all segments in a database record that can fit in one logical record. The index provides direct access to the root segment (and therefore to database records). The second data set, called the *overflow data set*, contains all segments in the database record that cannot fit in the primary data set. A key-sequenced data set (KSDS) is the primary data set and an entry-sequenced data set (ESDS) is the overflow data set.

There are several things you need to know about storage of HISAM database records:

- You define the logical record length of both the primary and overflow data set (subject to the rules listed in this chapter). The logical record length can be different for each data set. This allows you to define the logical record length in the primary data set as large enough to hold an “average” database record or the most frequently accessed segments in the database record. Logical record length in the overflow data set can then be defined (subject to some restrictions) as whatever is most efficient given the characteristics of your database records.
- Logical records are grouped into control intervals (CIs). A control interval is the unit of data transferred between an I/O device and storage. You define the size of CIs.
- Each database record starts at the beginning of a logical record in the primary data set. A database record can only occupy one logical record in the primary data set, but overflow segments of the database record can occupy more than one logical record in the overflow data set.
- Segments in a database record cannot be split and stored across two logical records. Because of this and because each database record starts a new logical record, unused space exists at the end of many logical records. When the database is initially loaded, IMS inserts a root segment with a key of all X'FF's as the last root segment in the database.

Figure 29 on page 71 shows four HISAM database records (shown in Figure 28) as they are initially stored on the primary and overflow data sets.

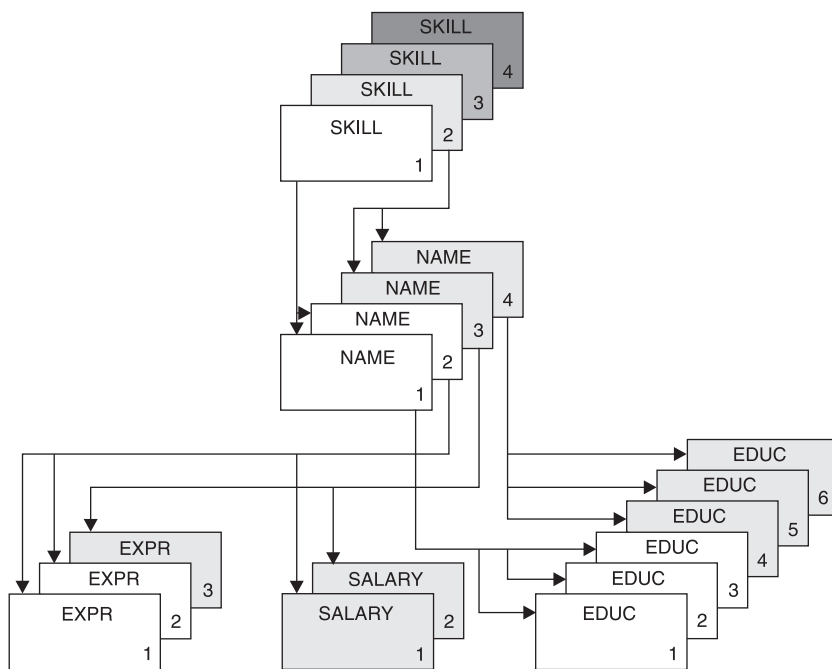


Figure 28. Example HISAM Database

In storage, a HISAM segment (see Figure 29) consists of a 2-byte prefix followed by user data. The first byte of the prefix is the segment code, which identifies the segment *type* to IMS. This number can be from 1 to 255. The segment code is assigned to the segment by IMS in ascending sequence, starting with the root segment and continuing through all dependents in hierarchic sequence. The second byte of the prefix is the delete byte.

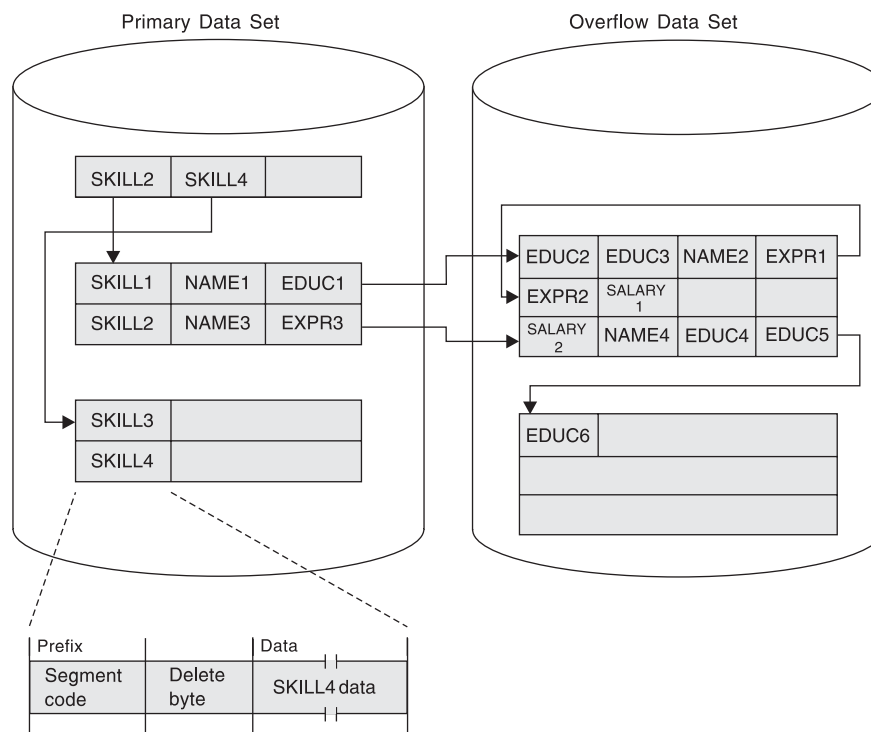


Figure 29. Example HISAM Database in Storage

Each logical record in the primary data set contains the root plus all dependents of the root (in hierarchic sequence) for which there is enough space. The remaining segments of the database record are put in the overflow data set (again in hierarchic sequence). The two “parts” of the database record are chained together with a direct-address pointer. When overflow segments in a database record use more than one logical record in the overflow data set (the case for the first and second database record in Figure 29), the logical records are also chained together with a direct-address pointer. Note in the figure that HISAM indexes do not contain a pointer to each root segment in the database. Rather, they point to the highest root key in each block or CI.

Diagnosis, Modification or Tuning Information

Figure 30 on page 72 illustrates the following points regarding the structure of a logical record in a HISAM database:

- In a logical record, the first 4 bytes are a direct-address pointer to the next logical record in the database record. This pointer maintains all logical records in a database record in correct sequence. The last logical record in a database record contains zeros in this field.
- Following the pointer are one or more segments of the database record in hierarchic sequence.
- Following the segments is a 1-byte segment code of 0. It says that the last segment in the logical record has been reached.




Figure 30. Format of a Logical Record in a HISAM Database

End of Diagnosis, Modification or Tuning Information

Accessing Segments

In HISAM, when an application program issues a call with a segment search argument (SSA) qualified on the key of the root segment, the segment is found by:

1. Searching the index for the first pointer with a value greater than or equal to the specified root key (the index points to the highest root key in each CI)
2. Following the index pointer to the correct CI
3. Searching this CI for the correct logical record (the root key value is compared with each root key in the CI)
4. When the correct logical record (and therefore database record) is found, searching sequentially through it for the specified segment

If an application program issues a GU call with an unqualified SSA for a root segment or with an SSA qualified on other than the root key, the HISAM index cannot be used. The search for the segment starts at the beginning of the database and proceeds sequentially until the specified segment is found.

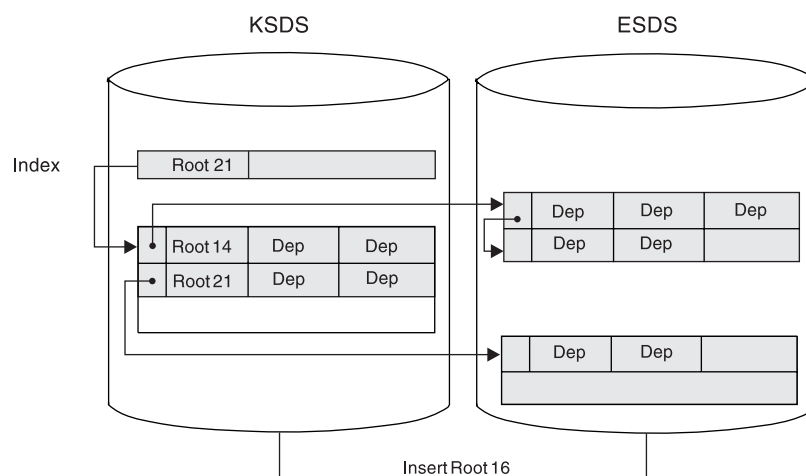
Inserting Root Segments Using VSAM

After an initial load, root segments inserted into a HISAM database are stored in the primary data set in ascending key sequence. The CI might or might not contain a free logical record into which the new root can be inserted. Both situations are described next.

A Free Logical Record Exists

Figure 31 on page 73 shows how insertion takes place when a free logical record exists. The new root is inserted into the CI in root key sequence. If there are logical records in the CI containing roots with higher keys, they are “pushed down” to create space for the new logical record.

BEFORE



AFTER

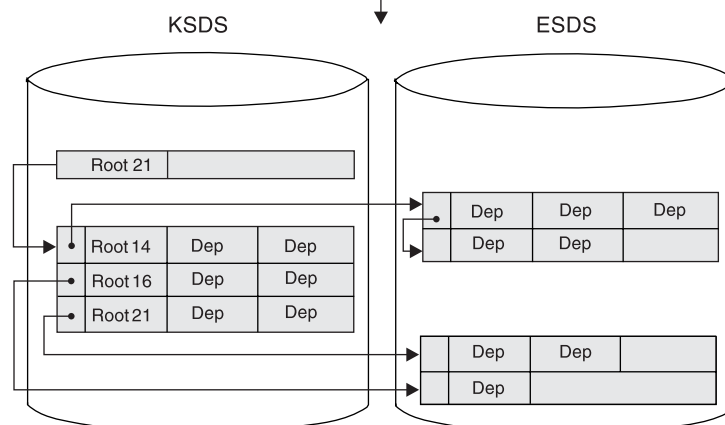


Figure 31. Inserting a Root Segment into a HISAM Database (Free Logical Record Exists in the CI)

No Free Logical Record Exists

Figure 32 on page 74 shows how insertion takes place when no free logical record exists in the CI. The CI is split forming two new CIs, both equal in size to the original one. Where the CI is split depends on what you have coded in the INSERT=parameter on the OPTIONS statement for the DFSVSAMP data set.

Related Reading: For information on the OPTIONS statement, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring* and Chapter 9, “Designing Full-Function Databases,” on page 245.

The split can occur at the point at which the root is inserted or midpoint in the CI. After the CI is split, free logical records exist in each new CI and the new root is inserted into the proper CI in root key sequence. If, as was the case in Figure 31, logical records in the new CI contained roots with higher keys, those logical records would be “pushed down” to create space for the new logical record.

When adding new root segments to a HISAM database, performance can be slightly improved if roots are added in ascending key sequence.

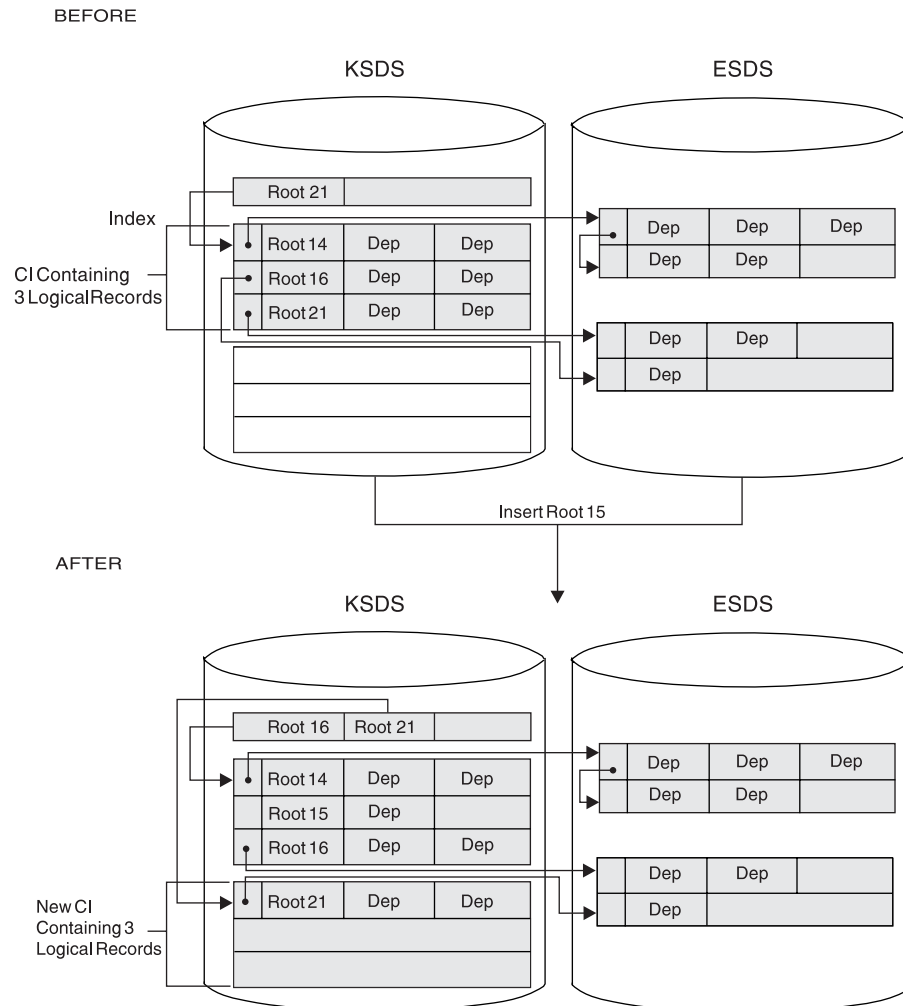


Figure 32. Inserting a Root Segment into a HISAM Database (No Free Logical Record Exists in the CI)

Inserting Dependent Segments

Dependent segments inserted into a HISAM database after initial load are inserted in hierarchic sequence. IMS decides where in the appropriate logical record the new dependent should be inserted. Two situations are possible. Either there is enough space in the logical record for the new dependent or there is not.

Figure 33 on page 75 shows how segment insertion takes place when there is enough space in the logical record. The new dependent is stored in its proper hierarchic position in the logical record by shifting the segments that hierarchically follow it to the right in the logical record.

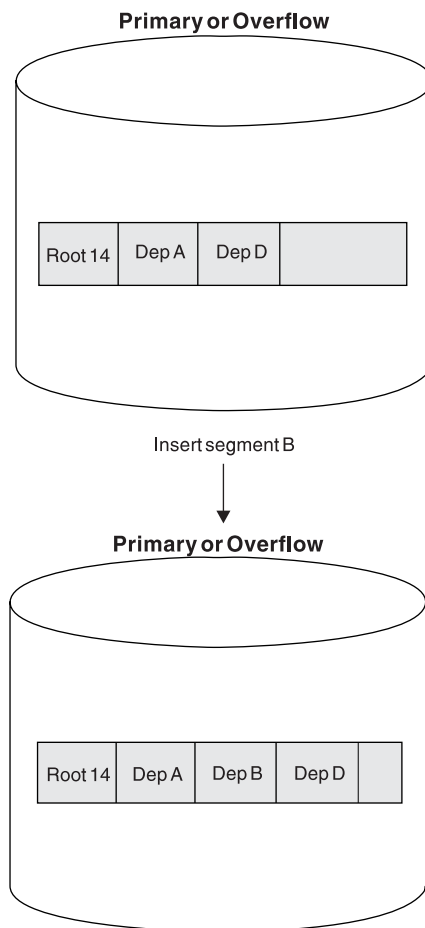
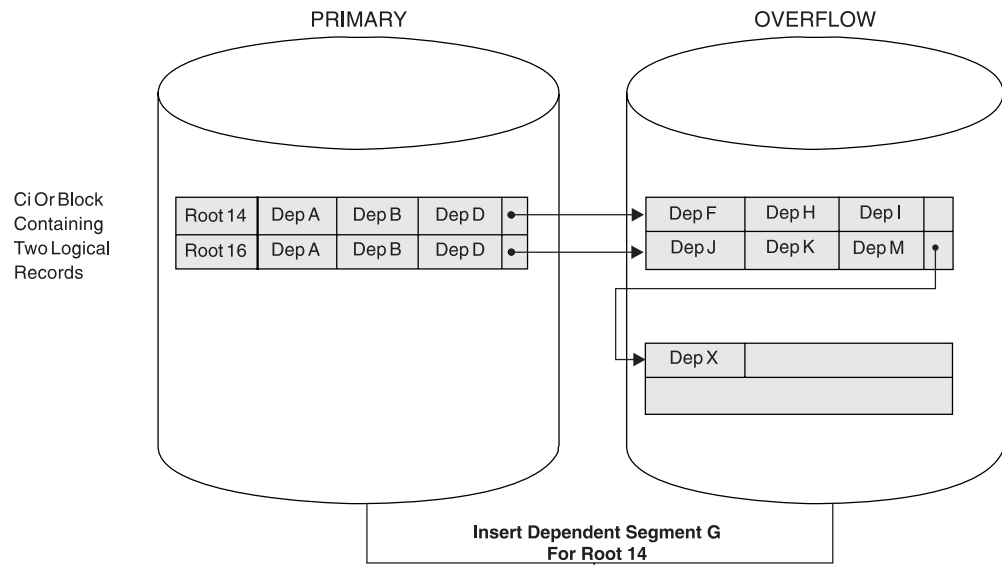


Figure 33. Inserting a Dependent Segment into a HISAM Database (Space Exists in the Logical Record)

Figure 34 on page 76 shows how segment insertion takes place when there is not enough space in the logical record. As in the previous case, new dependents are always stored in their proper hierarchic sequence in the logical record. However, all segments to the right of the new segment are moved to the first empty logical record in the overflow data set.

BEFORE



AFTER

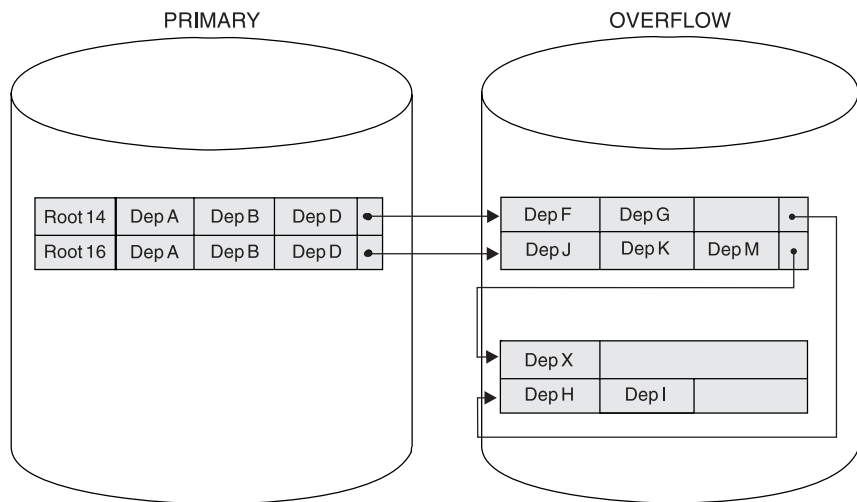


Figure 34. Inserting a Dependent Segment into a HISAM Database (No Space Exists in the Logical Record)

Deleting Segments

When segments are deleted from a HISAM database, they are marked as deleted in the delete byte in their prefix. They are not physically removed from the database; the one exception to this is discussed later in this topic. Dependent segments of the deleted segment are not marked as deleted, but because their parent is, the dependent segments cannot be accessed. These unmarked segments (as well as segments marked as deleted) are deleted when the database is reorganized.

One thing you should note is that when a segment is accessed that hierarchically follows deleted segments in a database record, the deleted segments must still be “searched through”. This concept is shown in Figure 35 and in Figure 36.

Segment B2 is deleted from this database record. This means that segment B2 and its dependents (C1, C2, and C3) can no longer be accessed, even though they still exist in the database.




Figure 35. The Hierarchic Segment Layout on the Database

A request to access segment D1 is made. Although segments B2, C1, C2, and C3 cannot be accessed, they still exist in the database. Therefore they must still be “searched through” even though they are inaccessible as shown in Figure 36.




Figure 36. Accessing a HISAM Segment That Hierarchically Follows Deleted Segments

In one situation, deleted segments are physically removed from the database. If the deleted segment is a root, the logical record containing the root is erased, provided neither the root nor any of its dependents is involved in a logical relationship. The default is ERASE=YES, and no “mark buffer altered” takes place. Thus a PROCOPT=G read job will not have to wait for locks after another job has set the delete byte, and will return a segment not found condition. To be consistent with other DB types, use ERASE=NO to cause a wait for physical delete prior to attempted read.

Related Reading: For more information on the ERASE parameter of the DBD statement, see the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

After the logical record is removed, its space is available for reuse. However, any overflow logical record containing dependents of this root is not available for reuse. Except for this special condition, you must unload and reload a HISAM database to regain space occupied by deleted segments.

Replacing Segments

Replacing segments in a HISAM database is straightforward as long as fixed length segments are being used. The data in the segment, once changed, is returned to its original location in storage. The key field in a segment cannot be changed.

The implications of replacing segments when variable-length segments are used is discussed under “Variable-Length Segments” on page 215.

Criteria for Selecting HISAM

You should use HISAM when you need sequential or direct access to roots and sequential processing of dependent segments in a database record. HISAM is a good choice of data organization when your database has most, or all, of the following characteristics.

- Each root has few dependents.

Root segment access is indexed, and is therefore fast. Dependent segment access is sequential, and is therefore slower.

- You have a small number of delete operations against the database.

Except for deleting root segments, all delete operations result in the creation of space that is unusable until the database is reorganized.

- Your applications depend on a small volume of root segments being inserted within a narrow key range (VSAM).

Root segments inserted after initial load are inserted in root key sequence in the appropriate CI in the KSDS. If many roots have keys within a narrow key range, many CI splits can occur. This will degrade performance.

- Most of your database records are about the same size.

The similar sizes allow you to pick logical record lengths and CI sizes so most database records fit on the primary data set. You want most database records to fit on the primary data set, because additional read and seek operations are required to access those parts of a database record on the overflow data set. Additional reads and seeks degrade performance. If, however, most of the processing you do against a database record occurs on segments in the primary data set (in other words, your high-use segments fit on the primary data set), these considerations might not be as important.

Having most of your database records the same size also saves space. Each database record starts at the beginning of a logical record. All space in the logical records not used by the database record is unusable. This is true of logical records in both the primary and overflow data set. If the size of your database records varies tremendously, large gaps of unused space can occur at the end of many logical records.

SHSAM, SHISAM and GSAM Databases

You typically use simple hierarchical sequential access method (SHSAM), simple hierarchical indexed sequential access method (SHISAM), and generalized sequential access method (GSAM) databases in two situations.

Situation 1 - Converting from a non-database system to IMS

SHSAM, SHISAM, and GSAM databases allow existing programs, using z/OS access methods, to remain usable during the conversion to IMS. This is possible because the format of the data in these databases is the same as in the z/OS data sets.

Situation 2 - Passing data

When a database (or non-database) application program passes data to a database (or non-database) application program, it first puts the data in a SHSAM, SHISAM, or GSAM database. The database (or non-database) application program then accesses the data from these databases.

The following topics describe each of the three database types:

- “SHSAM Databases”
- “SHISAM Databases”
- “GSAM Databases” on page 80

Table 8 on page 81 is a chart comparing SHSAM, SHISAM, and GSAM.

SHSAM Databases

A simple HSAM (SHSAM) database is an HSAM database containing only one type of segment, a root segment. The segment has no prefix, because no need exists for a segment code (there is only one segment type) or for a delete byte (deletes are not allowed).

SHSAM databases can be accessed by z/OS BSAM and QSAM because SHSAM segments contain user data only (no IMS prefixes). The ISRT, DLET, and REPL calls cannot be used to update. However, ISRT can be used to load an SHSAM database. Only GET calls are valid for processing an SHSAM database. These allow retrieval only of segments from the database. To update an SHSAM database, it must be reloaded. The situations in which SHSAM is typically used are explained in the introduction to this topic. Before deciding to use SHSAM, read the topic on GSAM databases, because GSAM has many of the same functions as SHSAM. Unlike SHSAM, however, GSAM files cannot be accessed from a message processing region. GSAM does allow you to take checkpoints and perform restart, though.

Although SHSAM databases can use the field-level sensitivity option, they *cannot* use any of the following options:

- Logical relationships
- Secondary indexing
- Multiple data set groups
- Variable-length segments
- Segment edit/compression exit routine
- Data Capture exit routines
- Logging, recovery, or reorganization

SHISAM Databases

A simple HISAM (SHISAM) database is a HISAM database containing only one type of segment, a root segment. The segment has no prefix, because no need exists for a segment code (there is only one segment type) or for a delete byte (deletes are done using a VSAM erase operation). SHISAM databases must be KSDSs; they are accessed through VSAM. Because SHISAM segments contain user data only (no IMS prefixes), they can be accessed by VSAM macros and DL/I calls. All the DL/I calls can be issued against SHISAM databases.

SHISAM IMS Symbolic Checkpoint Call

In addition to those situations described in the introduction to this topic, SHISAM is useful if you need an application program that accesses z/OS data sets to use the IMS symbolic checkpoint call.

The IMS symbolic checkpoint call makes restart easier than the z/OS basic checkpoint call. If the z/OS data set the application program is using is converted to a SHISAM database data set, the symbolic checkpoint call can be used. This allows application programs to take checkpoints during processing and then restart their programs from a checkpoint. The primary advantage of this is that, if the system fails, application programs can recover from a checkpoint rather than lose all processing that has been done. One exception applies to this: An application program for initially loading a database that uses VSAM as the operating system access method cannot be restarted from a checkpoint. Application programs using GSAM databases can also issue symbolic checkpoint calls. Application programs using SHSAM databases cannot.

Before deciding to use SHISAM, you should read the next topic on GSAM databases. GSAM has many of the same functions as SHISAM. Unlike SHISAM, however, GSAM files cannot be accessed from a message processing region.

SHISAM databases can use field-level sensitivity and Data Capture exit routines, but they *cannot* use any of the following options:

- Logical relationships
- Secondary indexing
- Multiple data set groups
- Variable-length segments
- Segment edit/compression exit routine

GSAM Databases

GSAM databases are sequentially organized databases designed to be compatible with z/OS data sets. GSAM databases can be on a data set previously created or one later accessed by the z/OS access methods VSAM or QSAM/BSAM. GSAM data sets can use fixed-length or variable-length records when VSAM is used, or fixed-length, variable-length or undefined-length records when QSAM/BSAM is used. If VSAM is used to process a GSAM database, the VSAM data set must be entry sequenced and on a DASD. If QSAM/BSAM is used, the physical sequential (DSORG=PS) data set can be placed on a DASD or tape unit. GSAM is designed to be compatible with z/OS data sets. The GSAM database has no hierarchy, database records, segments or keys.

GSAM IMS Symbolic Checkpoint Call

In addition to those situations described in the introduction to this topic, GSAM is useful if you need an application program that accesses z/OS data sets to use the IMS symbolic checkpoint call. The IMS symbolic checkpoint call makes restart easier than the z/OS basic checkpoint call. This IMS symbolic checkpoint call allows application programs to take checkpoints during processing, thereby allowing programs to restart from a checkpoint. A checkpoint call forces any GSAM buffers with inserted records to be written as short blocks. The primary advantage of taking checkpoints is that, if the system fails, the application programs can recover from a checkpoint rather than lose all your processed data. However, any application program that uses VSAM as an operating system access method and initially loads the database cannot be restarted from a checkpoint.

In general, always use DISP=OLD for GSAM data sets when restarting from a checkpoint even if you used DISP=MOD on the original execution of the job step. If you use DISP=OLD, the data set is positioned at its beginning. If you use DISP=MOD, the data set is positioned at its end.

Because GSAM databases are supported in a DCCTL environment, you may use them when you need to process sequential non-IMS data sets using a BMP program.

GSAM databases are loaded in the order in which you present records to the load program. You cannot issue DLET and REPL calls against GSAM databases; however, you can issue ISRT calls after the database is loaded but only to add records to the end of the data set. Records are not randomly added to a GSAM data set.

Although random processing of GSAM and SHSAM databases is possible, random processing of a GSAM database is done using a GU call qualified with a record search argument (RSA). This processing is primarily useful for establishing position in the database before issuing a series of GN calls.

Although SHSAM and SHISAM databases can be processed in any processing region, GSAM databases can only be processed in a batch or batch message processing region.

The following IMS options do not apply to GSAM databases:

- Logical relationships
- Secondary indexing
- Segment edit/compression exit routine
- Field-level sensitivity
- Data Capture exit routines
- Logging or reorganization
- Multiple data set groups

If you have application programs that need access to both IMS and z/OS data sets, you can use SHSAM, SHISAM, or GSAM. Which one you use depends on what functions you need. Table 8 compares the characteristics and functions available for each of the three types of databases.

Table 8. Comparison of SHSAM, SHISAM, and GSAM Databases

| Characteristics and Functions | SHSAM | SHISAM | GSAM |
|---|--------------|------------------|------------------|
| Hierarchic structure applicable? | NO | NO | NO |
| Segment prefix exist? | NO | NO | NO |
| Variable-length records used? | NO | NO | YES |
| Checkpoint/restart possible? | NO | YES ¹ | YES ¹ |
| Compatible with non-IMS data sets? | YES | YES | YES |
| Can VSAM be used as the operating system access method? | NO | YES | YES |
| Can BSAM be used as the operating system access method? | YES | NO | YES |
| Accessible from a batch region? | YES | YES | YES |

SHSAM, SHISAM, and GSAM Databases

Table 8. Comparison of SHSAM, SHISAM, and GSAM Databases (continued)

| Characteristics and Functions | SHSAM | SHISAM | GSAM |
|--|------------------|--------|------------------|
| Accessible from a batch message processing region? | YES | YES | YES |
| Accessible from a message processing region? | YES | YES | NO |
| Logging available? | NO | YES | NO |
| GET calls allowed? | YES | YES | YES |
| ISRT calls allowed? | YES ² | YES | YES ³ |
| Supported for CICS-DBCTL? | YES | YES | NO |
| Supported for DCCTL? | NO | NO | YES |

Note:

1. Using symbolic checkpoints
2. To load database only
3. Allowed only at the end of the data set

HDAM, PHDAM, HIDAM, and PHIDAM Databases

Hierarchical direct access method (HDAM) and hierarchical indexed direct access method (HIDAM) databases, which have many similarities, are referred to as HD databases. These HD databases can be partitioned using either the HALDB Partition Definition utility (DSPXPDDU) or DBRC commands and are then described as High Availability Large Databases (HALDBs). After you partition an HDAM database, it becomes a partitioned hierarchical direct access method (PHDAM) database. After you partition a HIDAM database, it becomes a partitioned hierarchical indexed direct access method (PHIDAM) database. Figure 37 illustrates a logical view of an HDAM and a PHDAM database.

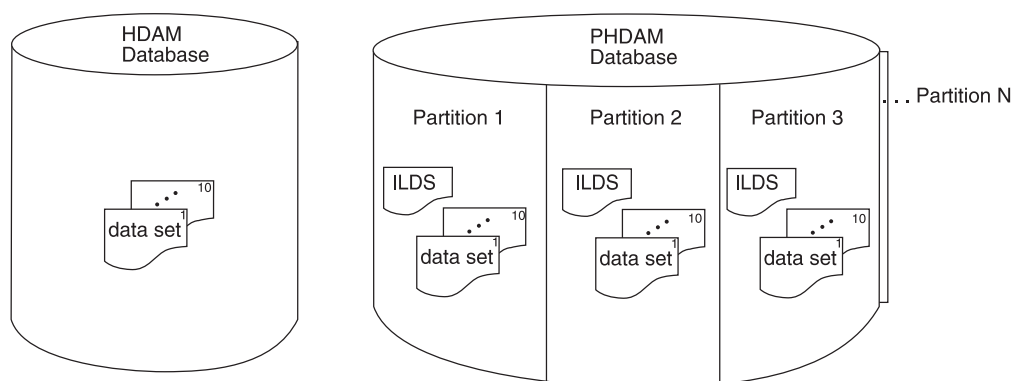


Figure 37. A Logical View of an HDAM and a PHDAM Database

HD databases differ from sequentially organized databases in two important ways. First, they use the direct method of storing data, and the hierarchic sequence of segments in the database is maintained by having segments point to one another. Except for a few special cases, each segment has one or more direct-address pointers in its prefix. When direct-address pointers are used, database records and segments can be stored anywhere in the database. Their position, once stored, is fixed, and they do not “move around” in the database when subsequent processing takes place. Instead, pointers are updated to reflect processing changes.

HD databases also differ from sequentially organized ones in that space in HD databases can be reused. If part or all of a database record is deleted, the deleted space can be reused when new database records or segments are inserted.

HD databases are stored on direct-access devices in either a VSAM ESDS or an OSAM data set. The storage organization in HDAM and HIDAM or PHDAM and PHIDAM is basically the same. Their primary difference is in the way their root segments are accessed. In HDAM or PHDAM, each root segment's storage location is found using a randomizing module. The randomizing module examines the root's key to determine the address of a pointer to the root segment. In HIDAM or PHIDAM, each root segment's storage location is found by searching an index. For HIDAM, this index is a database that IMS loads and maintains. The advantage of the HDAM randomizing module is that the I/O operations required to search an index are eliminated.

Figure 38 illustrates a logical view of a HIDAM and a PHIDAM database.

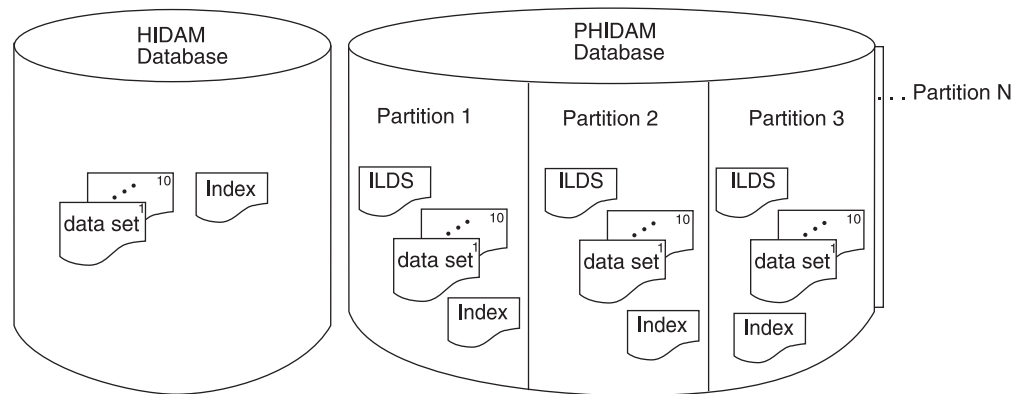


Figure 38. A Logical View of a HIDAM and a PHIDAM

Maximum Sizes of HD Databases

The maximum possible size of HDAM, PHDAM, HIDAM, and PHIDAM databases is based on the number of data sets the database can hold and the size of the data sets. The maximum possible size of a data set differs depending on whether VSAM or OSAM is used and whether the database is partitioned. Table 9 lists the maximum data set size, maximum number of data sets, and maximum database size for HDAM, PHDAM, HIDAM, and PHIDAM databases.

Table 9. Maximum Sizes for HDAM, HIDAM, PHDAM, and PHIDAM Databases

| Data Set Type | Maximum Data Set Size | Maximum Number of Data Sets | Maximum Database Size |
|-------------------------------|-----------------------|---|-----------------------|
| OSAM HDAM or HIDAM Database | 8 GB | 10 data sets | 80 GB |
| VSAM HDAM or HIDAM Database | 4 GB | 10 data sets | 40 GB |
| OSAM PHDAM or PHIDAM Database | 4 GB | 10 010 data sets (10 data sets per partition; 1001 partitions per database) | 40 040 GB |

HDAM, PHDAM, HIDAM, and PHIDAM

Table 9. Maximum Sizes for HDAM, HIDAM, PHDAM, and PHIDAM Databases (continued)

| Data Set Type | Maximum Data Set Size | Maximum Number of Data Sets | Maximum Database Size |
|-------------------------------|-----------------------|---|-----------------------|
| VSAM PHDAM or PHIDAM Database | 4 GB | 10 010 data sets (10 data sets per partition; 1001 partitions per database) | 40 040 GB |

Related Reading: For information on OSAM data sets, see Appendix C, “Using OSAM as the Access Method,” on page 513.

DL/I Calls Issuable Against HD Databases

All DL/I calls can be issued against HD databases. In addition, the following options are available:

- Multiple data set groups
- Logical relationships
- Secondary indexing
- Variable-length segments
- Segment edit/compression exit routine
- Data Capture exit routines
- Field-level sensitivity
- Logging, recovery, and offline reorganization
- Online reorganization for HALDB partitions

Related Reading:

- Except for logging and recovery, each of these options is discussed in detail in the topics of this chapter. For information on logging and recovery, see *IMS Version 9: Operations Guide*.
- For information on the online reorganization of HALDB partitions, see “HALDB Online Reorganization” on page 370.

When to Use HDAM and PHDAM

HDAM and PHDAM databases are typically used for direct access to database records. The randomizing module provides fast access to the root segment (and therefore the database record). HDAM and PHDAM databases also give you fast access to paths of segments as specified in the DBD in a database record. For example, in Figure 39 on page 85, if physical child pointers are used, they can be followed to reach segments B, C, D, or E. A hierarchic search of segments in the database record is bypassed. Segment B does not need to be accessed to get to segments C, D, or E. And segment D does not need to be accessed to get to segment E. Only segment A must be accessed to get to segment B or C. And only segments A and C must be accessed to get to segments D or E.

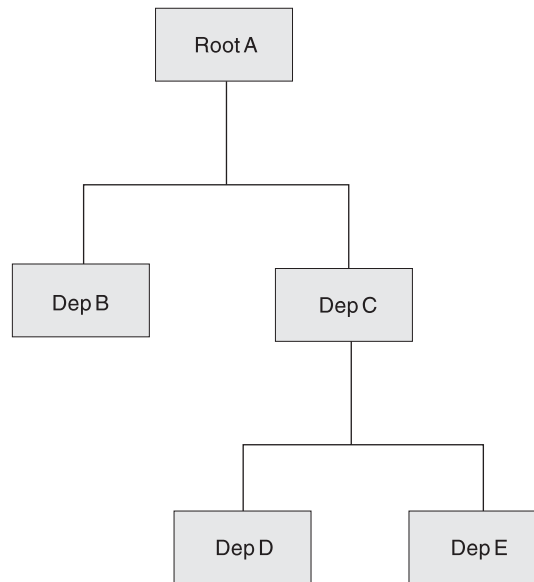


Figure 39. Example Database Record

When to Use HIDAM and PHIDAM

HIDAM and PHIDAM databases are typically used when you need both random and sequential access to database records and random access to paths of segment in a database record. Access to root segments (and therefore database records) is not as fast as with HDAM (or PHDAM), because the HIDAM (or PHIDAM) index database has to be searched for a root segment's address. However, because the index keeps the address of root segments stored in key sequence, database records can be processed sequentially.

What You Need to Know About HD Databases

Before looking in detail at how HD databases are stored and processed, you need to become familiar with:

- The various types of pointers you can specify for a HD database
- The general format of the database
- The use of special fields in the database

Types of Pointers You Can Specify

The hierarchic sequence of segments in a database record using the sequential access methods is maintained by keeping segments physically adjacent to each other in storage. In the HD access methods, segments in a database record are kept in hierarchic sequence using direct-address pointers. Except for a few special cases, each prefix in an HD segment contains one or more pointers. Each pointer is 4 bytes long and consists of the relative byte address of the segment to which it points. Relative, in this case, means relative to the beginning of the data set.

Several different types of direct-address pointers exist, and you will see how each works in the topics that follow in this section. However, there are three basic types:

- Hierarchic pointers, which point from one segment to the next in either forward or forward and backward hierarchic sequence

HDAM, PHDAM, HIDAM, and PHIDAM

- Physical child pointers, which point from a parent to each of its first or first and last children, for each child segment type
- Physical twin pointers, which point forward or forward and backward from one segment occurrence of a segment type to the next, under the same parent

When segments in a database record are typically processed in hierarchic sequence, use hierarchic pointers. When segments in a database record are typically processed randomly, use a combination of physical child and physical twin pointers. One thing to keep in mind while reading about pointers is that the different types, subject to some rules, can be mixed within a database record. However, because pointers are specified by segment type, all occurrences of the same segment type have the same type of pointer.

Each type of pointer is examined separately in this topic. The topic “Mixing Pointers” on page 93, discusses how pointers can be mixed. In the subtopics in this topic, each type of pointer is illustrated, and the database record on which each illustration is based is shown in Figure 40.

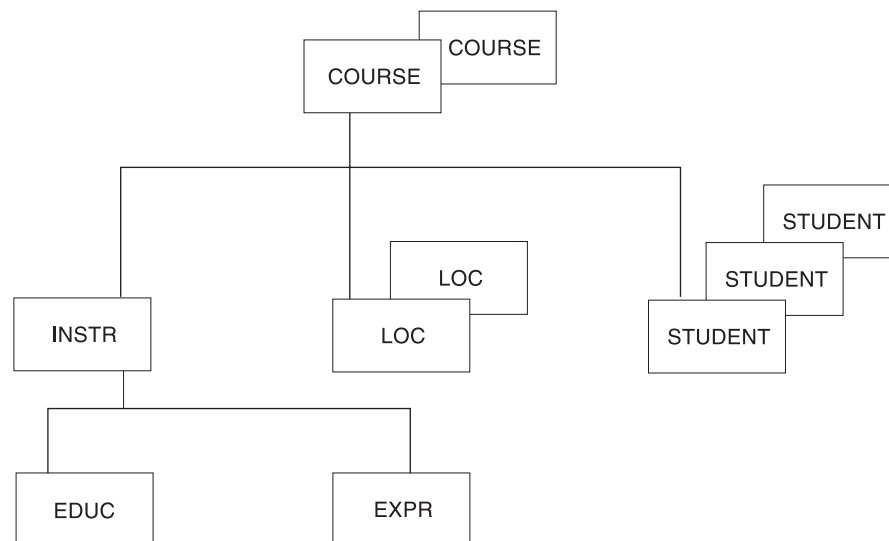


Figure 40. Example Database Record for Illustrating Pointers

Hierarchic Forward Pointers

With hierarchic forward (HF) pointers, each segment in a database record points to the segment that follows it in the hierarchy.

Figure 41 on page 87 shows hierarchic forward pointers:

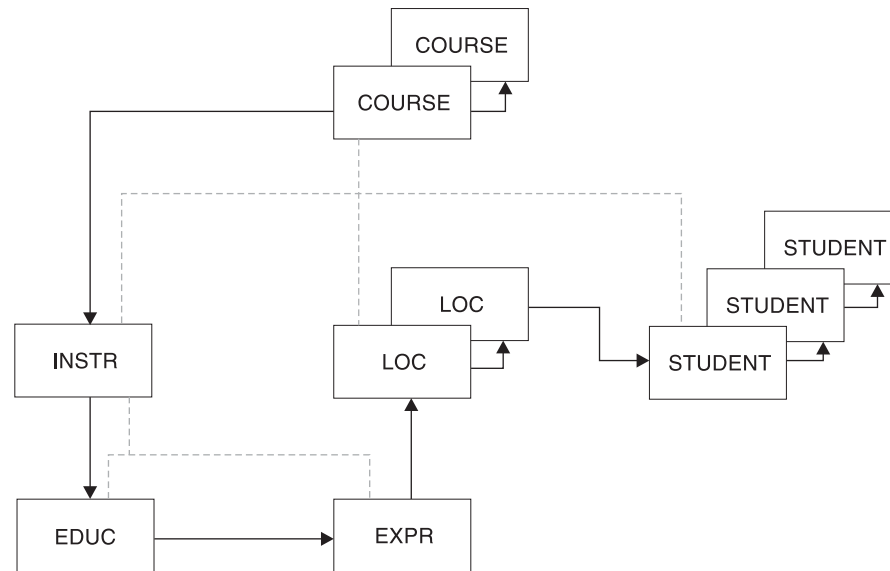


Figure 41. Hierarchic Forward Pointers

When an application program issues a call for a segment, HF pointers are followed until the specified segment is found. In this sense, the use of HF pointers in an HD database is similar to using a sequentially organized database. In both, to reach a dependent segment all segments that hierarchically precede it in the database record must be examined. HF pointers should be used when segments in a database record are typically processed in hierarchic sequence and processing does not require a significant number of delete operations. If there are a lot of delete operations, hierarchic forward and backward pointers (explained next) might be a better choice.

Four bytes are needed in each dependent segment's prefix for the HF pointer. Eight bytes are needed in the root segment. More bytes are needed in the root segment because the root points to both the next root segment and first dependent segment in the database record. HF pointers are specified by coding PTR=H in the SEGM statement in the DBD.

Restriction: HALDBs do not support HF pointers.

Hierarchic Forward and Backward Pointers

With hierarchic forward and backward pointers (HF and HB), each segment in a database record points to both the segment that follows and the one that precedes it in the hierarchy (except dependent segments do not point back to root segments). HF and HB pointers must be used together, since you cannot use HB pointers alone. Figure 42 on page 88 shows how HF and HB pointers work.

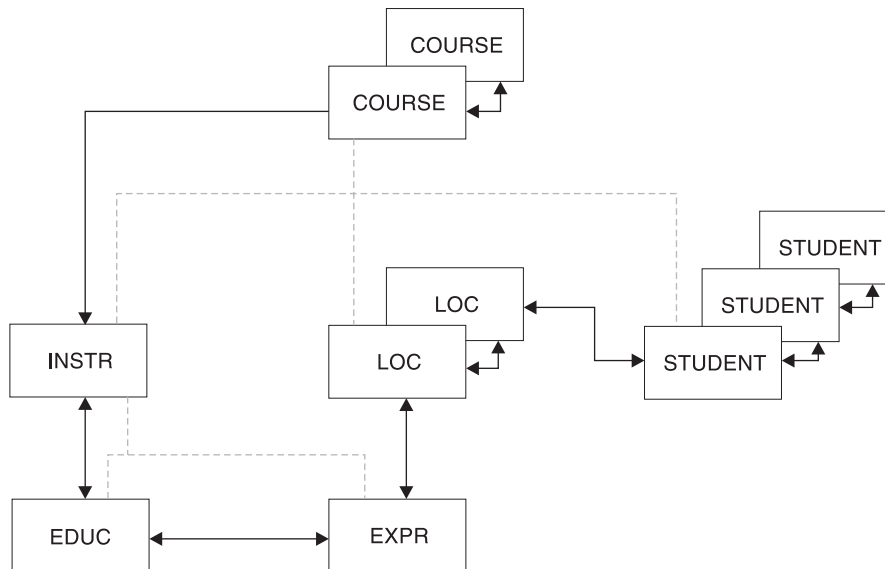


Figure 42. Hierarchic Forward and Backward Pointers

HF pointers work in the same way as the HF pointers described in “Hierarchic Forward Pointers” on page 86.

HB pointers point from a segment to one immediately preceding it in the hierarchy. In most cases, HB pointers are not required for delete processing. IMS saves the location of the previous segment retrieved on the chain and uses this information for delete processing. The backward pointers are useful for delete processing if the previous segment on the chain has not been accessed. This happens when the segment to be deleted is entered by a logical relationship.

The backward pointers are useful only when all of the following are true:

- Direct pointers from logical relationships or secondary indexes point to the segment being deleted or one of its dependent segments.
- These pointers are used to access the segment.
- The segment is deleted.

Eight bytes are needed in each dependent segment’s prefix to contain HF and HB pointers. Twelve bytes are needed in the root segment. More bytes are needed in the root segment because the root points:

- Forward to a dependent segment
- Forward to the next root segment in the database
- Backward to the preceding root segment in the database

HF and HB pointers are specified by coding PTR=HB in the SEGM statement in the DBD.

Restriction: HALDBs do not support HF and HB pointers.

Physical Child First Pointers

With physical child first (PCF) pointers, each parent segment in a database record points to the first occurrence of each of its immediately dependent child segment types. Figure 43 shows PCF pointers:

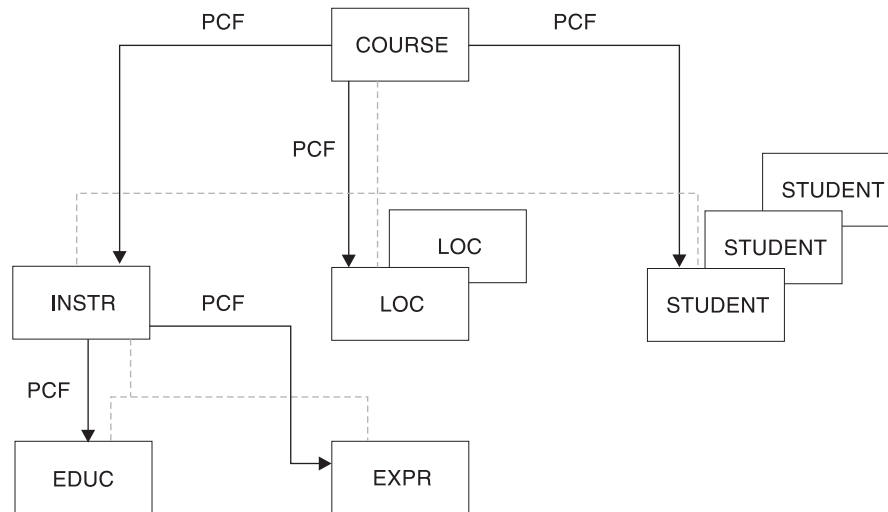


Figure 43. Physical Child First Pointers

With PCF pointers, the hierarchy is only partly connected. No pointers exist to connect occurrences of the same segment type under a parent. Physical twin pointers (explained in “Types of Pointers You Can Specify” on page 85) can be used to form this connection. Use PCF pointers when segments in a database record are typically processed randomly and either sequence fields are defined for the segment type, or if not defined, the insert rule is FIRST or HERE. If sequence fields are not defined and new segments are inserted at the end of existing segment occurrences, the combination of PCF and physical child last (PCL) pointers (explained next) can be a better choice.

Related Reading:

- For more information on insert rules, see *IMS Version 9: Application Programming: Database Manager*.
- For information on specifying insert rules using the RULES= parameter of the SEGM segment definition statement, see *IMS Version 9: Utilities Reference: System*.

Four bytes are needed in each parent segment for each PCF pointer. PCF pointers are specified by coding PARENT=((name,SNGL)) in the SEGM statement in the DBD. This is the SEGM statement for the child being pointed to, not the SEGM statement for the parent. Note, however, that the pointer is stored in the parent segment.

Physical Child First and Last Pointers

With physical child first and last pointers (PCF and PCL), each parent segment in a database record points to both the first and last occurrence of its immediately dependent child segment types. PCF and PCL pointers must be used together, since you cannot use PCL pointers alone. Figure 44 shows PCF and PCL pointers:

HDAM, PHDAM, HIDAM, and PHIDAM

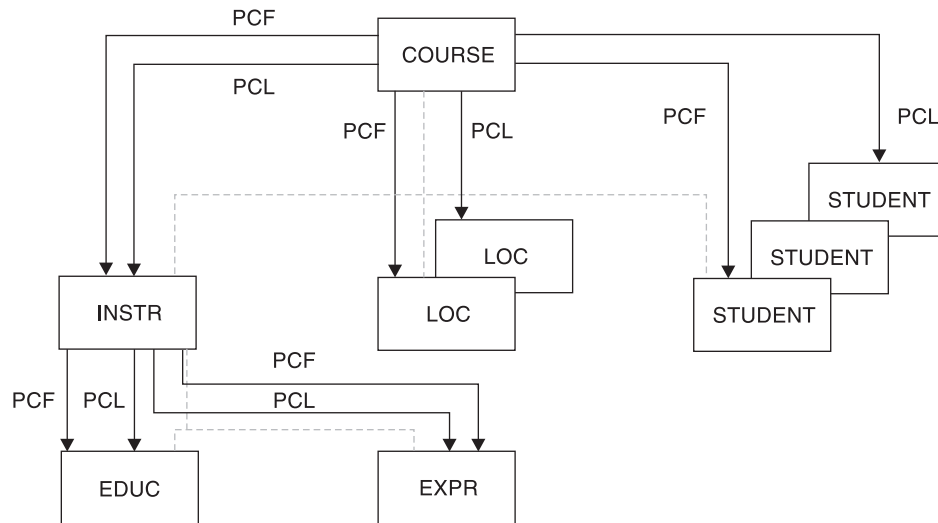


Figure 44. Physical Child First and Last Pointers

Note that if only one physical child of a particular parent segment exists, the PCF and PCL pointers both point to the same segment. As with PCF pointers, PCF and PCL pointers leave the hierarchy only partly connected, and no pointers exist to connect occurrences of the same segment type under a parent. Physical twin pointers (explained in “Types of Pointers You Can Specify” on page 85) can be used to form this connection.

PCF and PCL pointers (as opposed to just PCF pointers) are typically used when:

- No sequence field is defined for the segment type.
- New segment occurrences of a segment type are inserted at the end of all existing segment occurrences.

On insert operations, if the ISRT rule of LAST has been specified, segments are inserted at the end of all existing segment occurrences for that segment type. When PCL pointers are used, fast access to the place where the segment will be inserted is possible. This is because there is no need to search forward through all segment occurrences stored before the last occurrence. PCL pointers also give application programs fast retrieval of the last segment in a chain of segment occurrences. Application programs can issue calls to retrieve the last segment by using an unqualified SSA with the command code L. When a PCL pointer is followed to get the last segment occurrence, any further movement in the database is forward.

A PCL pointer does not enable you to search from the last to the first occurrence of a series of dependent child segment occurrences.

Four bytes are needed in each parent segment for each PCF and PCL pointer. PCF and PCL pointers are specified by coding the PARENT= operand in the SEGM statement in the DBD as PARENT=((name,DBLE)). This is the SEGM statement for the child being pointed to, not the SEGM statement for the parent. Note, however, that the pointers are stored in the parent segment.

A parent segment can have SNGL specified on one immediately dependent child segment type and DBLE specified on another.

Figure 45 on page 91 shows the result of specifying PCF and PCL pointers in the following DBD.


```

DBD
SEGM A
SEGM B  PARENT=((name.SNGL))  (specifies PCF pointer only)
SEGM C  PARENT=((name.DBLE))  (specified PCF and PCL pointers)

```




Figure 45. Specifying PCF and PCL Pointers

Physical Twin Forward Pointers

With physical twin forward (PTF) pointers, each segment occurrence of a given segment type under the same parent points forward to the next segment occurrence. Figure 46 on page 92 illustrates this.

Note that, except in PHIDAM databases, PTF pointers can be specified for root segments. When this is done in an HDAM or PHDAM database, the root segment points to the next root in the database chained off the same root anchor points (RAP). If no more root segments are chained from this RAP, the PTF pointer is zero.

Related Reading: For more information on RAPs, see “General Format of HD Databases and Use of Special Fields” on page 95.

When PTF pointers are specified for root segments in a HIDAM database, the root segment does *not* point to the next root in the database. For an explanation of where the root segment points, see “Use of RAPs in a HIDAM Database” on page 102.

If you specify PTF pointers on a root segment in a HIDAM database, the HIDAM index must be used for all sequential processing of root segments. Using only PTF pointers increases access time. You can eliminate this overhead by specifying PTF and physical twin backward (PTB) pointers (discussed in “Physical Twin Forward and Backward Pointers” on page 92).

You cannot use PTF pointers for root segments in a PHIDAM database. PHIDAM databases only support PTF pointers for dependent segments.

With PTF pointers, the hierarchy is only partly connected. No pointers exist to connect parent and child segments. Physical child pointers can be used to form this connection. PTF pointers should be used when segments in a database record are typically processed randomly, and you do not need sequential processing of database records.

Four bytes are needed for the PTF pointer in each segment occurrence of a given segment type. PTF pointers are specified by coding PTR=T in the SEGM statement in the DBD. This is the SEGM statement for the segment containing the pointer. The combination of PCF and PTF pointers is used as the default when pointers are not specified in the DBD. Figure 46 show PTF pointers:

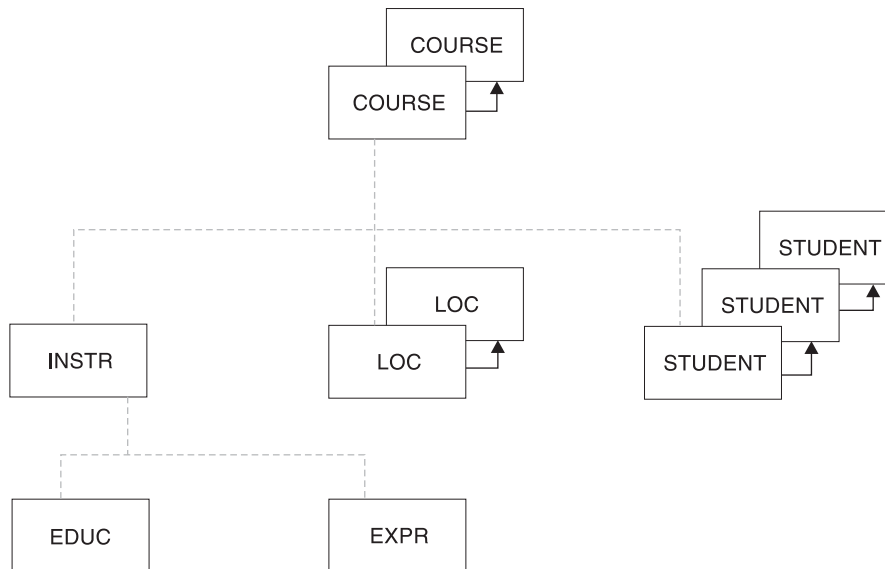


Figure 46. Physical Twin Forward Pointers

Physical Twin Forward and Backward Pointers

With physical twin forward and backward (PTF and PTB) pointers, each segment occurrence of a given segment type under the same parent points both forward to the next segment occurrence and backward to the previous segment occurrence. PTF and PTB pointers must be used together, since you cannot use PTB pointers alone. Figure 47 illustrates how PTF and PTB pointers work.

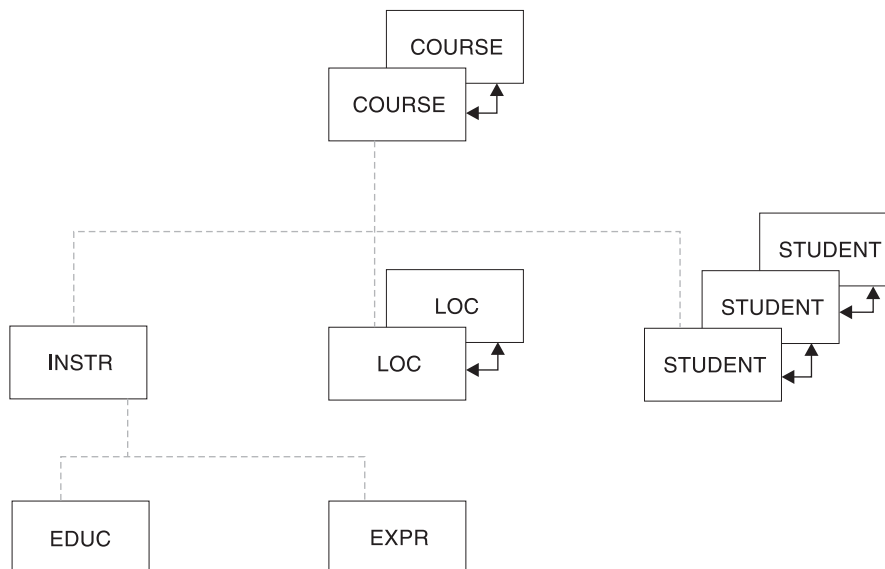


Figure 47. Physical Twin Forward and Backward Pointers

Note that PTF and PTB pointers can be specified for root segments. When this is done, the root segment points to both the next and the previous root segment in the database. As with PTF pointers, PTF and PTB pointers leave the hierarchy only partly connected. No pointers exist to connect parent and child segments. Physical child pointers (explained previously) can be used to form this connection.

PTF and PTB pointers (as opposed to just PTF pointers) should be used on the root segment of a HIDAM or a PHIDAM database when you need fast sequential processing of database records. By using PTB pointers in root segments, an application program can sequentially process database records without IMS' having to refer to the HIDAM or PHIDAM index. For HIDAM databases, PTB pointers improve performance when deleting a segment in a twin chain accessed by a virtually paired logical relationship. Such twin-chain access occurs when a delete from the logical access path causes DASD space to be released.

Eight bytes are needed for the PTF and PTB pointers in each segment occurrence of a given segment type. PTF and PTB pointers are specified by coding PTR=TB in the SEGM statement in the DBD.

Mixing Pointers

Because pointers are specified by segment type, the various types of pointers can be mixed within a database record. However, only hierarchic or physical, but not both, can be specified for a given segment type. The types of pointers that can be specified for a segment type are:

| | |
|--------------------|------------------------------------|
| HF | Hierarchic forward |
| HF and HB | Hierarchic forward and backward |
| PCF | Physical child first |
| PCF and PCL | Physical child first and last |
| PTF | Physical twin forward |
| PTF and PTB | Physical twin forward and backward |

Figure 48 on page 94 shows a database record in which pointers have been mixed. Note that, in some cases, for example, dependent segment B, many pointers exist even though only one type of pointer is or can be specified. Also note that if a segment is the last segment in a chain, its last pointer field is set to zero (the case for segment E1, for instance). One exception is noted in the rules for mixing pointers. Figure 48 has a legend that explains what specification in the PTR= or PARENT= operand causes a particular pointer to be generated.

The rules for mixing pointers are:

- If PTR=H is specified for a segment, no PCF pointers can exist from that segment to its children. For a segment to have PCF pointers to its children, you must specify PTR=T or TB for the segment.
- If PTR=H or PTR=HB is specified for the root segment, the first child will determine if an H or HB pointer is used. All other children must be of the same type.
- If PTR=H is specified for a segment other than the root, PTR=TB and PTR=HB cannot be specified for any of its children. If PTR=HB is specified for a segment other than the root, PTR=T and PTR=H cannot be specified for any of its children.

That is, the child of a segment that uses hierarchic pointers must contain the same number of pointers (twin or hierarchic) as the parent segment.

- If PTR=T or TB is specified for a segment whose immediate parent used PTR=H or PTR=HB, the last segment in the chain of twins does not contain a zero. Instead, it points to the first occurrence of the segment type to its right on the same level in the hierarchy of the database record. This is true even if no twin *chain* yet exists, just a single segment for which PTR=T or TB is specified (dependent segment B and E2 in the figure illustrate this rule).

HDAM, PHDAM, HIDAM, and PHIDAM

- If PTR=H or HB is specified for a segment whose immediate parent used PTR=T or TB, the last segment in the chain of twins contains a zero (dependent segment C2 in the figure illustrates this rule).

Figure 48 shows an example of mixing pointers in a database record.

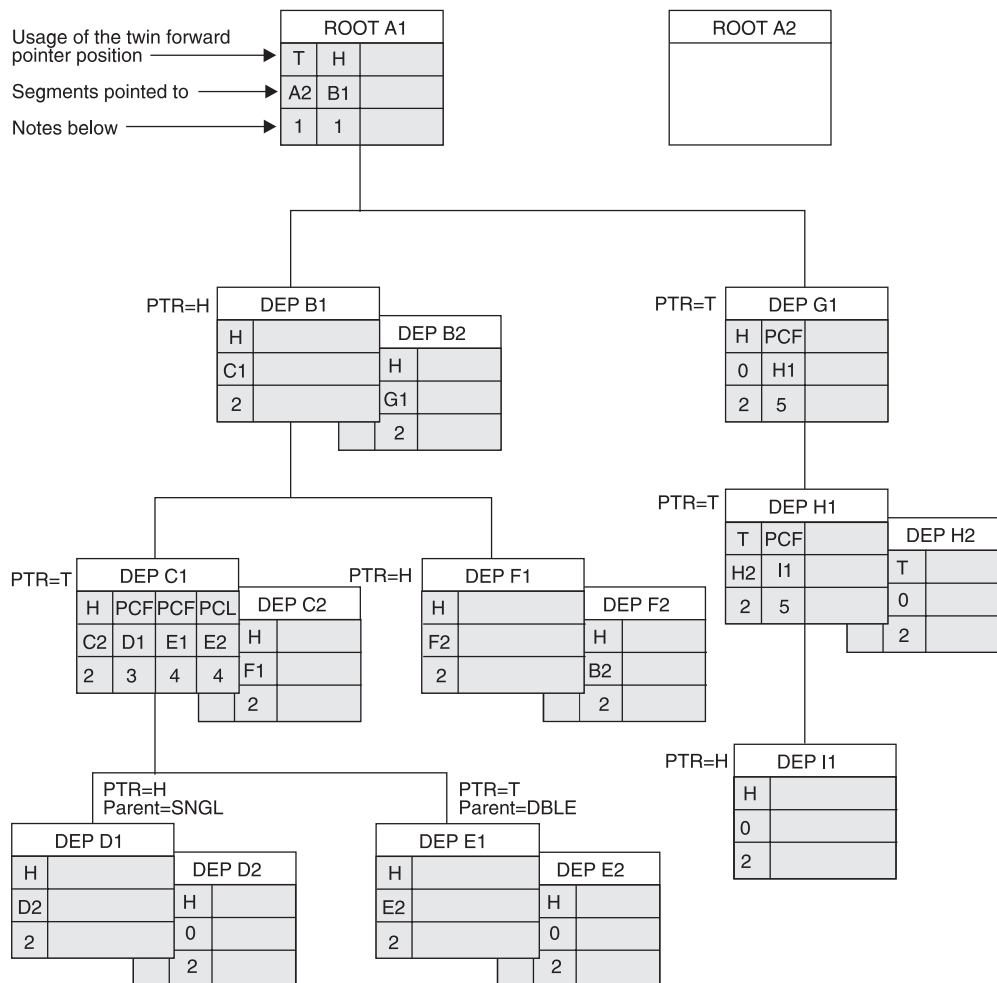


Figure 48. Mixing Pointers

Notes for Figure:

1. These pointers are generated when you specify PTR=H on the root segment.
2. If you specify PTR=H, usage is hierarchical (H); otherwise usage is twin (T).
3. These pointers are generated when you specify PTR=T on segment type C and PARENT=SNGL on segment type D
4. These pointers are generated when you specify PTR=T on segment type C and PARENT=DBLE on segment type E
5. These pointers are generated when you specify PTR=T on this segment type

Sequence of Pointers in a Segment's Prefix

When a segment contains more than one type of pointer, pointers are put in the segment's prefix in the following sequence:

1. HF
2. HB

Or:

1. PF
2. PTB
3. PCF
4. PCL

General Format of HD Databases and Use of Special Fields

The way in which an HD database is organized is not particularly complex, but some of the special fields in the database used for things like managing space make HD databases seem quite different from sequentially organized databases. This topic looks at the general layout of the database special fields.

The databases referred to here are the HDAM or PHDAM and the HIDAM or PHIDAM databases. HIDAM and PHIDAM each have an additional database, the primary index database; for HIDAM, you allocate it; for PHIDAM, IMS allocates it; for both, IMS maintains the index. This topic examines the index database when dealing with the storage of HIDAM records. Figure 49 shows the general format of an HD database and some of the special fields used in it.

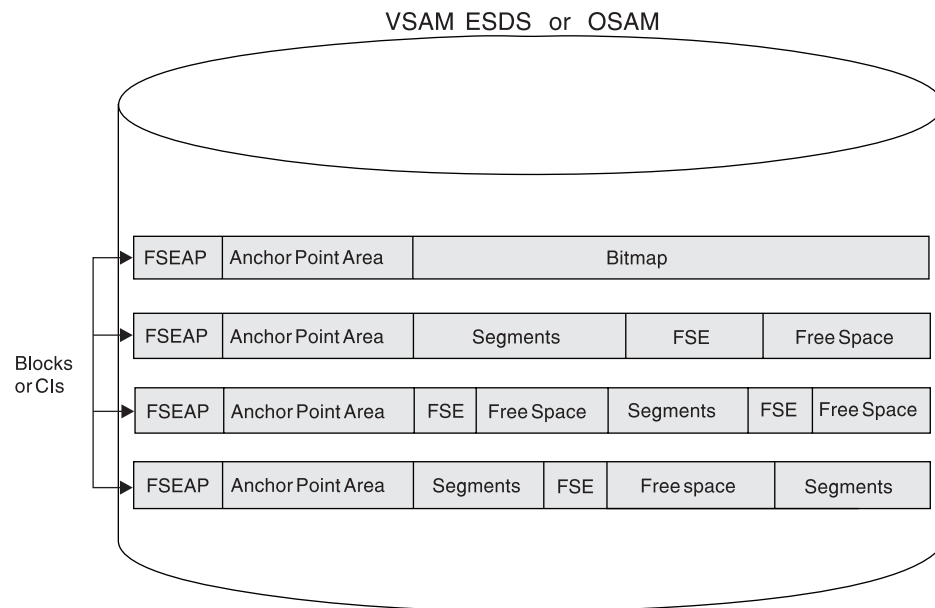


Figure 49. Format of an HD Database and Special Fields in It

HD databases use a single data set, that is either a VSAM ESDS or an OSAM data set. The data set contains one or more CIs (VSAM ESDS) or blocks (OSAM). Database records in the data set are in unblocked format. Logical record length is the same as the block size when OSAM is used. When VSAM is used, logical record length is slightly less than CI size. (VSAM requires some extra control information in the CI.) You can either specify logical record length yourself or have it done by the Database Description Generation (DBDGEN) utility. The utility generates logical record lengths equal to a quarter, third, half, or full track block.

All segments in HD Databases begin on a halfword boundary. If a segment's total length is an odd number, the space used in an HD database will be one byte longer than the segment. The extra byte is called a "slack byte".

HDAM, PHDAM, HIDAM, and PHIDAM

Note that the database in Figure 49 contains areas of free space. This free space could be the result of delete or replace operations done on segments in the data set. Remember, space can be reused in HD databases. Or it could be free space you set aside when loading the database. HD databases allow you to set aside free space by specifying that periodic blocks or CIs be left free or by specifying that a percentage of space in each block or CI be left free.

Examine the four fields illustrated in Figure 49. Three of the fields are used to manage space in the database. The remaining one, the anchor point area, contains the addresses of root segments. The fields are:

- *Bit map*. Bit maps contain a string of bits. Each bit describes whether enough space is available in a particular CI or block to hold an occurrence of the longest segment defined in the data set group. The first bit says whether the CI or block that the bit map is in has free space. Each consecutive bit says whether the next consecutive CI or block has free space. When the bit value is one, it means the CI or block has enough space to store an occurrence of the longest segment type you have defined in the data set group. When the bit value is zero, not enough space is available.

The first bit map in an OSAM data set is in the first block of the first extent of the data set. In VSAM data sets, the second CI is used for the bit map and the first CI is reserved. The first bit map in a data set contains n bits that describe space availability in the next $n-1$ consecutive CIs or blocks in the data set. After the first bit map, another bit map is stored at every n th CI or block to describe whether space is available in the next group of CIs or blocks in the data set.

For a HALDB partition, the first bit map block stores the partition ID (2 bytes) and the reorganization number (2 bytes). These are stored before the FSEAP at the beginning of the block.

An example bit map is shown in Figure 50.

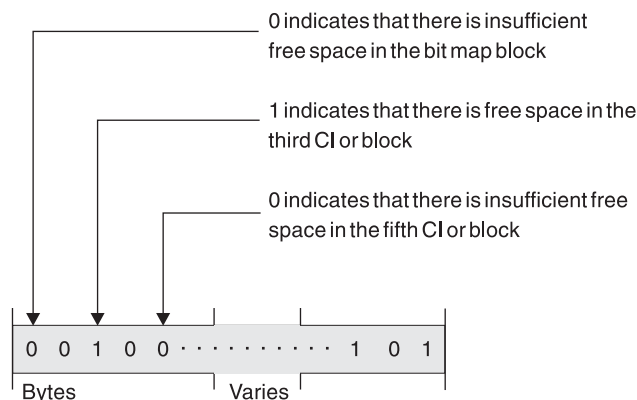


Figure 50. Bit Map for HD Databases

- *Free space element anchor point (FSEAP)*. FSEAPs are made up of two 2-byte fields. The first contains the offset, in bytes, to the first free space element (FSE) in the CI or block. FSEs describe areas of free space in a block or CI. The second field identifies whether this block or CI contains a bit map. If the block or CI does not contain a bit map, the field is zeros. One FSEAP exists at the beginning of every CI or block in the data set. IMS automatically generates and maintains FSEAPs.

An FSEAP is shown in Figure 51 on page 97.

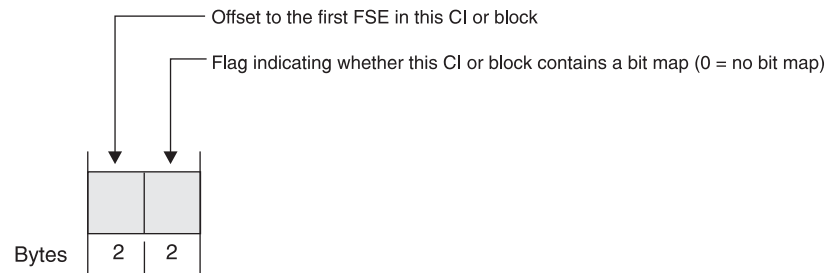


Figure 51. An FSEAP

The FSEAP in the first bit map block in an OSAM data set has a special use. It is used to contain the DBRC usage indicator for the database. The DBRC usage indicator is used at database open time for update processing to verify usage of the correct DBRC RECON data set.

- **Free space element (FSE).** An FSE describes each area of free space in a CI or block that is 8 or more bytes in length. IMS automatically generates and maintains FSEs. FSEs occupy the first 8 bytes of the area that is free space. FSEs consist of three fields:
 - **Free space chain pointer (CP) field.** This field contains, in bytes, the offset from the beginning of this CI or block to the next FSE in the CI or block. This field is 2 bytes long. The CP field is set to zero if this is the last FSE in the block or CI.
 - **Available length (AL) field.** This field contains, in bytes, the length of the free space identified by this FSE. The value in this field includes the length of the FSE itself. The AL field is 2 bytes long.
 - **Task ID (ID) field.** This field contains the task ID of the program that freed the space identified by the FSE. The task ID allows a given program to free and reuse the same space during a given scheduling without contending for that space with other programs. The ID field is 4 bytes long.

An FSE is shown in Figure 52.

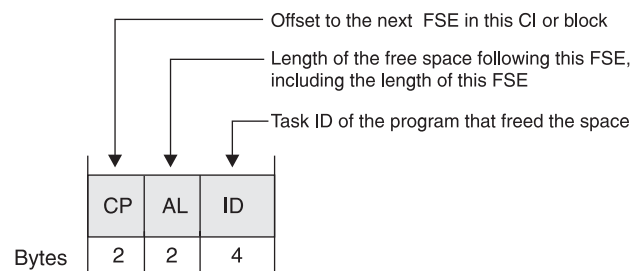


Figure 52. An FSE

- **Anchor point area.** The anchor point area is made up of one or more 4-byte root anchor points (RAPs). Each RAP contains the address of a root segment. For HDAM, you specify the number of RAPs you need on the RMNAME parameter in the DBD statement. For PHDAM, you specify the number of RAPs you need on the RMNAME parameter in the DBD statement, or by using the HALDB Partition Definition utility, or on the DBRC INIT.PART command. For HIDAM (but not PHIDAM), you specify whether RAPs exist by specifying PTR=T or PTR=H for a root segment type. Only one RAP per block or CI is generated. How RAPs are used in HDAM, PHDAM, and HIDAM differs. Therefore RAPs will be examined further in the following topics:

HDAM, PHDAM, HIDAM, and PHIDAM

- “How HDAM and PHDAM Records Are Stored”
- “How HIDAM and PHIDAM Records Are Stored” on page 100

An anchor point area in an HDAM or PHDAM database is shown in Figure 53.




Figure 53. An HDAM or PHDAM Anchor Point Area

How HDAM and PHDAM Records Are Stored

HDAM or PHDAM databases consist of two parts: a root addressable area and an overflow area. The root addressable area contains root segments and is the primary storage area for dependent segments in a database record. The overflow area is for the storage of segments that do not fit in the root addressable area. You specify the size of the root addressable area in the relative block number (RBN) operand of the RMNAME parameter in the DBD statement. For PHDAM, you can also use the HALDB Partition Definition utility to specify the size of the root addressable area. You also specify the maximum number of bytes of a database record to be stored in the root addressable area by using the BYTES operand of the RMNAME parameter in the DBD statement. For PHDAM databases, you can use the HALDB Partition Definition utility to specify the maximum number of bytes in the root addressable area.

Figure 54 shows sample Skills database records. Figure 55 on page 99 shows how these records are stored in a HDAM or HIDAM database.

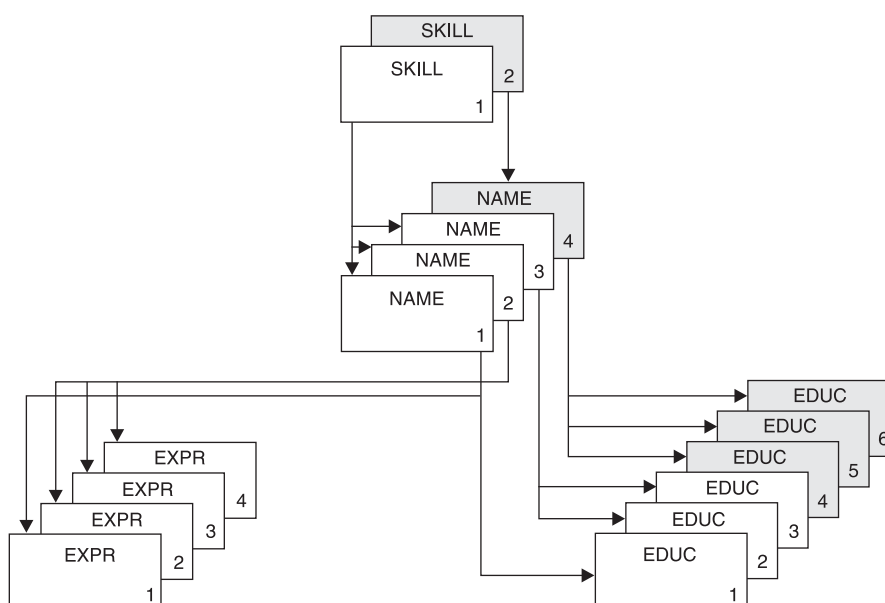


Figure 54. Two Example HD Database Records

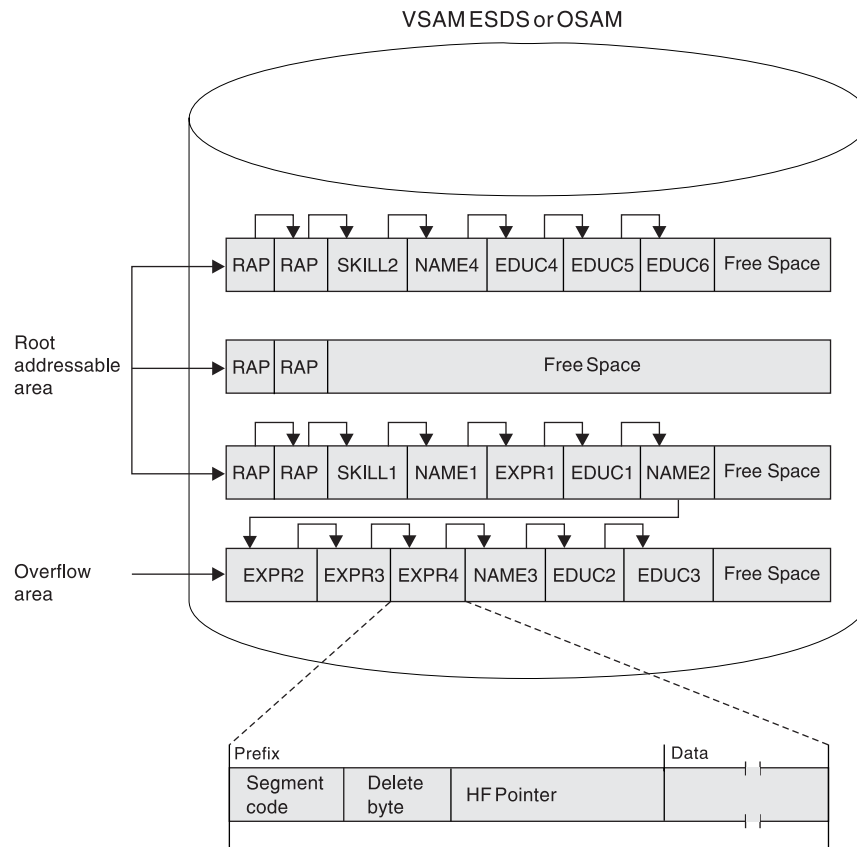


Figure 55. HDAM or PHDAM Database Records in Storage

When the database is initially loaded, the root and each dependent segment are put in the root addressable area until the next segment to be stored will cause the total space used to exceed the amount of space you specified in the BYTES operand. At this point, all remaining dependent segments in the database record are stored in the overflow area.

In an HDAM or a PHDAM database, the order in which you load database records does not matter. The user randomizing module determines where each root is stored. However, as with all types of databases, when the database is loaded, all dependents of a root must be loaded in hierarchic sequence following the root.

To store an HDAM or a PHDAM database record, the randomizing module takes the root's key and, by hashing or some other arithmetic technique, computes an RBN or CI number and a RAP number within the block or CI. The module gives these numbers to IMS, and IMS determines where in the root addressable area to store the root. The RBN or CI tells IMS in which CI or block (relative to the beginning of the data set) the RAP will be stored. The RAP number tells which RAP in the CI or block will contain the address of the root. During load, IMS stores the root and as many of its dependent segments that will fit (based on the bytes operand) in the root addressable area.

When the database is initially loaded, it puts the root and segments in the first available space in the specified CI or block, if this is possible. IMS then puts the 4-byte address of the root in the RAP of the CI or block designated by the randomizing module. RAPs only exist in the root addressable area. If space is not

HDAM, PHDAM, HIDAM, and PHIDAM

available in the root addressable area for a root, it is put in the overflow area. The root, however, is chained from a RAP in the root addressable area.

When Not Enough Root Storage Room Exists

If the CI or block specified by the randomizing module does not contain enough room to store the root, IMS uses the HD space search algorithm to find space. This algorithm is explained in “How the HD Space Search Algorithm Works” on page 107. When insufficient space exists in the specified CI or block to store the root, the algorithm finds the closest available space to the specified CI or block. When space is found, the address of the root is still stored in the specified RAP in the original block or CI generated by the randomizing module.

If the randomizing module generates the same relative block and RAP number for more than one root, the RAP points to a single root and all additional roots with the same relative block and RAP number are chained to each other using physical twin pointers. Roots are always chained in ascending key sequence. If non-unique keys exist, the ISRT rules of FIRST, LAST, and HERE determine the sequence in which roots are chained (These ISRT rules are explained in *IMS Version 9: Application Programming: Database Manager*). All roots chained like this from a single anchor point area are called *synonyms*.

Figure 55 on page 99 shows two HDAM or PHDAM database records and how they appear in storage after initial load. In this example, enough space exists in the specified block or CI to store the roots, and the unique relative block and RAP numbers for each root generated by the randomizing module. The bytes parameter specifies enough space for five segments of the database record to fit in the root addressable area. All remaining segments are put in the overflow area. When HDAM or PHDAM database records are initially loaded, dependent segments that cannot fit in the root addressable area are simply put in the first available space in the overflow area.

Note how segments in the database record are chained together. In this case, hierarchic pointers are used instead of the combination of physical child/physical twin pointers. Each segment points to the next segment in hierarchic sequence. Also note that two RAPs were specified per CI or block and each of the roots loaded is pointed to by a RAP. For simplicity, Figure 55 on page 99 does not show the various space management fields.

An HDAM or PHDAM segment in storage (see Figure 55 on page 99) consists of a prefix followed by user data. The first byte of the prefix is the segment code, which identifies the segment type to IMS. This number can be from 1 to 255. The segment code is assigned to the segment type by IMS in ascending sequence, starting with the root segment and continuing through all dependents in hierarchic sequence. The second byte of the prefix is the delete byte. The third field in the prefix contains the one or more addresses of segments to which this segment is pointing. In this example, hierarchic forward pointers are used. Therefore, the EXPR4 segment contains only one address, the address of the NAME3 segment.

How HIDAM and PHIDAM Records Are Stored

A HIDAM database is actually composed of two databases. One database contains the database records and the other database contains the HIDAM index. HIDAM uses the index to get to a specific root segment rather than the root anchor points that HDAM and PHDAM use.

Loading a HIDAM or PHIDAM Database

Root segments in a HIDAM or PHIDAM database must have a unique key field, because an index entry exists for each root segment based on the root's key. When initially loading a HIDAM or a PHIDAM database, you should present all root segments to the load program in ascending key sequence, with all dependents of a root following in hierarchic sequence. Figure 56 shows how the two Skills database records shown in Figure 54 on page 98 appear in storage after initial load. Note that HIDAM or PHIDAM, unlike HDAM or PHDAM, have no root addressable or overflow area, just a series of blocks or CIs. When database records are initially loaded, they are simply loaded one after another in the order in which they are presented to the load program. The space in Figure 56 at the end of each block or CI is free space specified when the database was loaded. In this example, 30% free space per block or CI was specified.

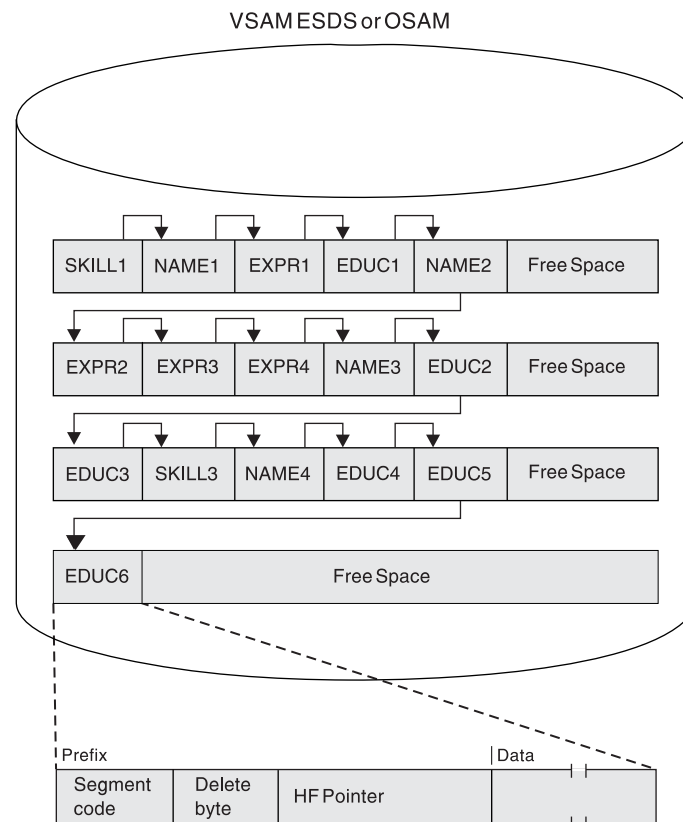


Figure 56. HIDAM or PHIDAM Database Records in Storage

Note how segments in a database record are chained together. In this case, hierarchic pointers were used instead of the combination of physical child/physical twin pointers. Each segment points to the next segment in hierarchic sequence. No RAPs exist in Figure 56. Although HIDAM databases can have RAPs, you probably do not need to use them. The reason for not using RAPs is explained in “Use of RAPs in a HIDAM Database” on page 102.

In storage, a HIDAM or PHIDAM segment (see Figure 56) consists of a prefix followed by user data. The first byte of the prefix is the segment code, which identifies the segment type to IMS. This number can be from 1 to 255. The segment code is assigned to the segment by IMS in ascending sequence, starting with the root segment and continuing through all dependents in hierarchic sequence. The second byte of the prefix is the delete byte. The third field in the prefix contains the

one or more addresses of segments to which this segment is pointing. In this example, hierarchic forward pointers are used. The EDUC6 segment contains only one address, the address of the root segment of the next database record (not shown here) in the database.

Creating an Index Segment

As each root is stored in a HIDAM or PHIDAM database, IMS creates an index segment for the root and stores it in the index database or data set. The index database consists of a VSAM KSDS. The KSDS contains an index segment for each root in the database or HALDB partition. When initially loading a HIDAM or PHIDAM database, IMS will insert a root segment with a key of all X'FF's as the last root in the database or partition.

The format of an index segment is shown in Figure 57.




Figure 57. Format of an Index Segment

The prefix portion of the index segment contains the delete byte and the root's address. The data portion of the index segment contains the key field of the root being indexed. This key field identifies which root segment the index segment is for and remains the reason why root segments in a HIDAM or PHIDAM database must have unique sequence fields. Each index segment is a separate logical record. Figure 58 shows the index database that IMS would generate when the two database records in Figure 56 on page 101 were loaded.

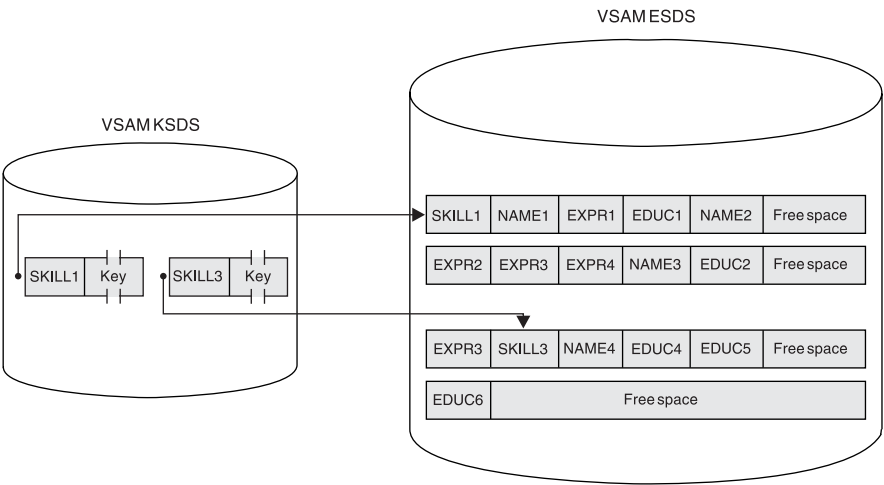


Figure 58. HIDAM or PHIDAM Index Databases

Use of RAPs in a HIDAM Database

RAPs are used differently in HIDAM databases than they are in HDAM or PHDAM databases. In HDAM or PHDAM, RAPs exist to point to root segments. When the randomizing module generates roots with the same relative block and RAP number (synonyms), the RAP points to one root and synonyms are chained together off that root.

In HIDAM databases, RAPs are generated only if you specify PTR=T or PTR=H for a root segment. When either of these is specified, one RAP is put at the beginning of each CI or block, and root segments within the CI or block are chained from the RAP in reverse order based on the time they were inserted. By this method, the RAP points to the last root inserted into the block or CI, and the hierarchic or twin forward pointer in the first root inserted into the block or CI is set to zero. The hierarchic or twin forward pointer in each of the other root segments in the block points to the previous root inserted in the block. Figure 59 shows what happens if you specify PTR=T or PTR=H for root segments in a HIDAM database.

Figure 59 uses the following abbreviations:

- FSE** Free space element
- RAP** Root anchor point
- SC** Segment code
- DB** Delete byte
- TF** Twin forward
- H** Hierarchic forward




Figure 59. Specifying PTR=T or PTR=H for Root Segments in a HIDAM Database

Note that if you specify PTR=H for a PHIDAM root, you get an additional hierarchic pointer to the first dependent in the hierarchy. In Figure 59, a “1” indicates where this additional hierarchic pointer would appear.

The implication of using PTR=T or PTR=H is that the pointer from one root to the next cannot be used to process roots sequentially. Instead, the HIDAM index must be used for all sequential root processing, and this increases access time. Specify PTR=TB or PTR=HB for root segments in a HIDAM database. Then no RAP is generated, and GN calls against root segments proceed along the normal physical twin forward chain. If no pointers are specified for HIDAM root segments, the default is PTR=T.

Accessing Segments

The way in which a segment in an HD database is accessed depends on whether the DL/I call for the segment is qualified or unqualified.

Qualified Calls

When a call is issued for a root segment and the call is qualified on the root segment's key, the way in which the database record containing the segment is found depends on whether the database is HDAM, PHDAM, HIDAM, or PHIDAM. In an HDAM or a PHDAM database, the randomizing module generates the root

HDAM, PHDAM, HIDAM, and PHIDAM

segment's (and therefore the database record's) location. In a HIDAM or a PHIDAM database, the HIDAM or PHIDAM index is searched until the index segment containing the root's key is found.

Once the root segment is found, if the qualified call is for a dependent segment, IMS searches for the dependent by following the pointers in each dependent segment's prefix. The exact way in which the search proceeds depends on the type of pointers you are using. Figure 60 shows how a dependent segment is found when PCF and PTF pointers are used.

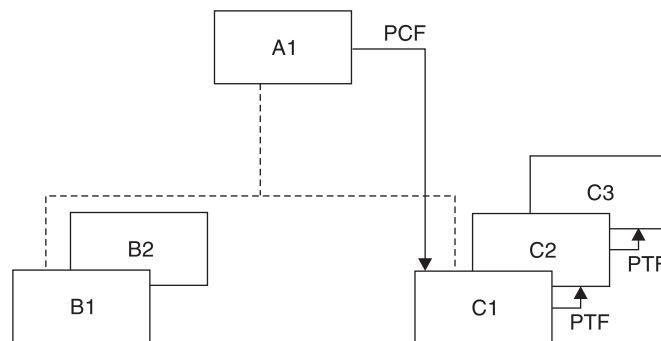


Figure 60. How Dependent Segments Are Found Using PCF and PTF Pointers

Unqualified Calls

When an unqualified call is issued for a segment, the way in which the search proceeds depends on:

- Whether the database is HDAM, PHDAM, HIDAM, or PHIDAM
- Whether a root or dependent segment is being accessed
- Where position in the database is currently established
- What type of pointers are being used
- Where parentage is set (if the call is a GNP)

Because of the many variables, it is not practical to generalize on how a segment is accessed.

Inserting Root Segments

The way in which a root segment is inserted into an HD database depends on whether the database is HDAM, PHDAM, HIDAM, or PHIDAM. For PHDAM or PHIDAM databases, partition selection is first performed based on the key of the root segment.

Inserting Root Segments into an HDAM or PHDAM Database

After initial load, root segments are inserted into an HDAM or PHDAM database in exactly the same way they are inserted during initial load. This process is explained in "How HDAM and PHDAM Records Are Stored" on page 98.

Inserting Root Segments Into a HIDAM or PHIDAM Database

After initial load, root segments are inserted into a HIDAM or PHIDAM database as follows (see Figure 61 on page 105):

1. The HIDAM or PHIDAM index is searched for an index segment with a root key greater than the key of the root to be inserted.
2. The new index segment is inserted in ascending root sequence.

- Once the index segment is created, the root segment is stored in the database at the location specified by the HD space search algorithm. How this algorithm works is described in “How the HD Space Search Algorithm Works” on page 107.

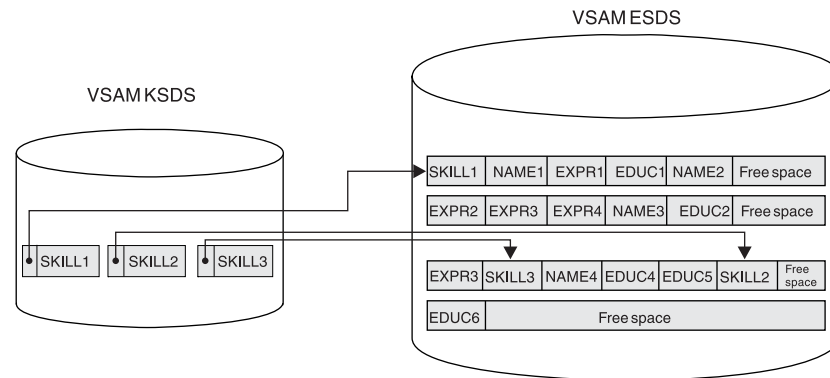


Figure 61. Inserting a Root Segment into a HIDAM or PHIDAM Database

Updating the Space Management Fields When a Root Segment Is Inserted

When a root segment is inserted into an HD database, the space management fields need to be updated. Figure 62 on page 106 illustrates this process. The figure makes several assumptions so real values could be put in the space management fields. These assumptions are:

- The database is HDAM or PHDAM (and therefore has a root addressable area).
- VSAM is the access method; therefore there are CIs (not blocks) in the database. Because VSAM is used, each logical record has 7 bytes of control information.
- Logical records are 512 bytes long.
- One RAP exists in each CI.
- The root segment to be inserted (SKILL1) is 32 bytes long.

The “before” picture shows the CI containing the bit map (in VSAM, the bit map is always in the second CI in the database). The second bit in the bit map is set to 1, which says there is free space in the next CI. In the next CI (CI #3):

- The FSEAP says there is an FSE (which describes an area of free space) 8 bytes from the beginning of this CI.
- The anchor point area (which has one RAP in this case) contains zeros because no root segments are currently stored in this CI.
- The FSE AL field says there is 497 bytes of free space available starting at the beginning of this FSE.

The SKILL1 root segment to be inserted is only 32 bytes long; therefore CI #3 has plenty of space to store SKILL1.

The “after” picture shows how the space management fields in CI #3 are updated when SKILL1 is inserted.

- The FSEAP now says there is an FSE 40 bytes from the beginning of this CI.
- The RAP points to SKILL1. The pointer value in the RAP is derived using the following formula:

$$\text{Pointer value} = (\text{CI size}) * (\text{CI number} - 1) + \text{Offset with the CI root segment}$$

HDAM, PHDAM, HIDAM, and PHIDAM

In this case, the pointer value is 1032 (pointer value = $512 \times 2 + 8$).

- The FSE has been “moved” to the beginning of the remaining area of free space. The FSE AL field says there is 465 bytes ($497 - 32$) of free space available, starting at the beginning of this FSE.

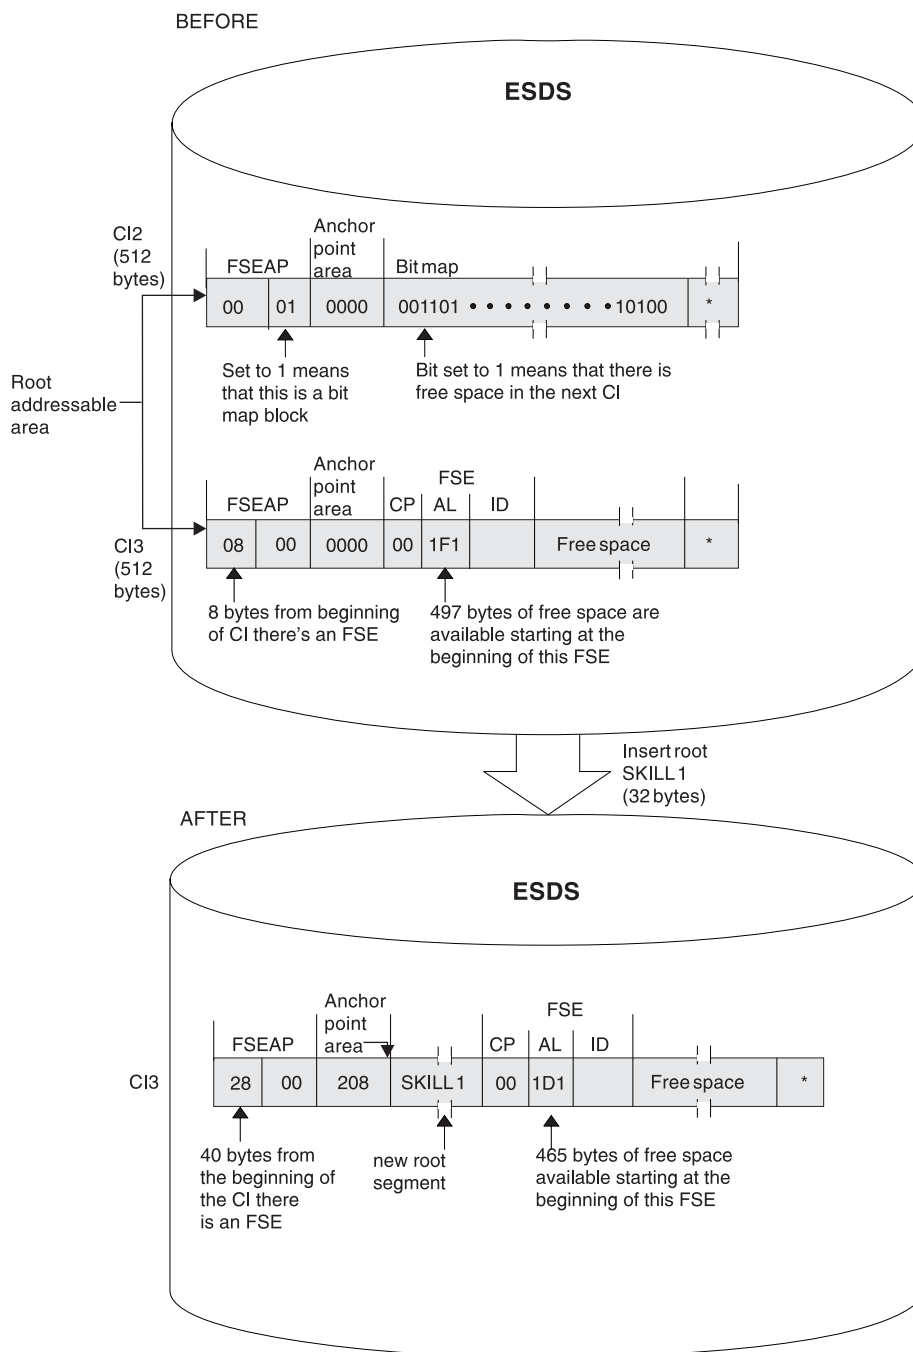


Figure 62. Updating the Space Management Fields in an HDAM or PHDAM Database

Inserting Dependent Segments

After initial load, dependent segments are inserted into HD databases using the HD space search algorithm. How this algorithm works is described in “How the HD Space Search Algorithm Works” on page 107.

As with the insertion of root segments into an HD database, the various space management fields in the database need to be updated (This process was explained and illustrated in “Updating the Space Management Fields When a Root Segment Is Inserted” on page 105).

Deleting Segments

When a segment is deleted in an HD database, it is physically removed from the database. The space it occupied can be reused when new segments are inserted. As with the insertion of segments into an HD database, the various space management fields need to be updated (This process was explained and illustrated in “Updating the Space Management Fields When a Root Segment Is Inserted” on page 105).

- The bit map needs to be updated if the block or CI from which the segment is deleted now contains enough space for a segment to be inserted. (Remember, the bit map says whether enough space exists in the block or CI to hold a segment of the longest type defined. Thus, if the deleted segment did not free up enough space for the longest segment type defined, the bit map is not changed.)
- The FSEAP needs to be updated to show where the first FSE in the block or CI is now located.
- When a segment is deleted, a new FSE might be created or the AL field value in the FSE that immediately precedes the deleted segment might need to be updated.
- If the deleted segment is a root segment in an HDAM or a PHDAM database, the value in its PTF pointer is put in the RAP or in the PTF pointer that pointed to it. This maintains the chain off the RAP and removes the deleted segment from the chain.

If a deleted segment is next to an already available area of space, the two areas are combined into one unless they are created by an online task that has not yet reached a sync point.

Replacing Segments

Replacing segments in HD databases is straightforward as long as fixed-length segments are used. The segment data, once changed, is simply returned to its original location in storage. The key field in a segment cannot be replaced.

Provided sufficient adjacent space is available, the segment data is returned to its original location when a variable-length segment is replaced with a longer segment. If adjacent space is unavailable, space is obtained from the overflow area for the lengthened data portion of the segment. This segment is referred to as a “separated data segment.” It has a 2-byte prefix consisting of a 1-byte segment code and a 1-byte delete flag, followed by the segment data. The delete byte of the separated data segment is set to X'FF', indicating that this is a separated data segment. A pointer is built immediately following the original segment to point to the separated data. Bit 4 of the delete byte of the original segment is set ON to indicate that the data for this segment is separated. The unused remaining space in the original segment is available for reuse.

How the HD Space Search Algorithm Works

The general rule for inserting a segment into an HD database is to store the segment (whether root or dependent) in the most desirable block or CI.

Root Segment

The most desirable block depends on the access method. For HDAM or PHDAM roots, the most desirable block is the one containing either the RAP or root segment that will point to the root being inserted. For HIDAM or PHIDAM roots, if the root does not have a twin backward pointer, the most desirable block is the one containing the root with the next higher key. If the root has a twin backward pointer, the most desirable block is the root with the next lower key.

Dependent Segment

The most desirable block is the one containing the segment that points to the inserted segment. If both physical child and physical twin pointers are used, the most desirable block is the one containing either the parent or the immediately-preceding twin. If hierarchic pointers are used, the most desirable block is the one containing the immediately-preceding segment in the hierarchy.

Second-Most Desirable Block

When it is not possible to store one or more segments in the most desirable block (space is not available), the HD space search algorithm searches for the second-most desirable block or CI. (This search is done only if the block is in the buffer pool or contains free space according to the bit map). The second-most desirable block or CI is a block or CI that was left free when the database was loaded or reorganized. Every *n*th block or CI can be left free by specifying the FRSPC= keyword in the DATASET macro of the DBDGEN utility. If you do not specify in the FRSPC= keyword that every *n*th block or CI be left free, the HD space search algorithm will not search for the second-most desirable block or CI.

Related Reading: For more information on the FRSPC= and SEARCHA= keywords, see *IMS Version 9: Utilities Reference: System*.

All search ranges defined in the HD space search algorithm, excluding steps 9 through 11, are limited to the physical extent that includes the most desirable block. When the most desirable block is in the overflow area, the search ranges, excluding steps 9 through 11, are restricted to the overflow area.

The steps in the HD space search algorithm follow. They are arranged in the sequence in which they are performed. The first time any one of the steps in the list results in available space, the search is ended and the segment is stored.

Look for space:

1. In the most desirable block (this block or CI is in the buffer pool).
2. In the second-most desirable block or CI.
3. In any block or CI in the buffer pool on the same cylinder.
4. In any block or CI on the same track, as determined by consulting the bit map. (The bit map says whether space is available for the longest segment type defined.)
5. In any block or CI on the same cylinder, as determined by consulting the bit map.
6. In any block or CI in the buffer pool within plus or minus *n* cylinders. Specify *n* in the SCAN= keyword in the DATASET statement in the DBD.
7. In any block or CI plus or minus *n* cylinders, as determined by consulting the bit map.
8. In any block or CI in the buffer pool at the end of the data set.

9. In any block or CI at the end of the data set, as determined by consulting the bit map. The data sets will be extended as far as possible before going to the next step.
10. In any block or CI in the data set where space exists, as determined by consulting the bit map. (This step is not used when a HIDAM or PHIDAM database is loaded.)

Some of the above steps are skipped in load mode processing.

If the dependent segment being inserted is at the highest level in a secondary data set group, the place and the way in which space is found differ:

- First, if the segment has no twins, steps 1 through 8 in the HD space search algorithm are skipped.
- Second, if the segment has a twin that precedes it in the twin chain, the most desirable block is the one containing that twin.
- Third, if the segment has only twins that follow it in the twin chain, the most desirable block is the one containing the twin to which the new segment is chained.

Locking Protocols

IMS uses locks to isolate the database changes made by concurrently executing programs. Locking is accomplished by using either the Program Isolation (PI) lock manager or the IRLM. The PI lock manager provides only four locking levels and the IRLM supports eleven lock states.

The IRLM also provides support for “feedback only” and “test” locking required, and it supplies feedback on lock requests compatible with feedback supplied by the PI lock manager.

Locking to Provide Program Isolation

For all database organizations, the basic item locked is the database record. The database record is locked when position is first obtained in it. The item locked is the root segment, or for HDAM or PHDAM, the anchor point. Therefore, for HDAM or PHDAM, all database records chained from the anchor are locked. The processing option of the PCB determines whether or not two programs can concurrently access the same database record. If the processing option permits updates, then no other program can concurrently access the database record. The database record is locked until position is changed to a different database record or until the program reaches a commit point.

When a program updates a segment with an INSERT, DELETE, or REPLACE call, the segment, not the database record, is locked. On an INSERT or DELETE call, at least one other segment is altered and locked.

Because data is always accessed hierarchically, when a lock on a root (or anchor) is obtained, IMS determines if any programs hold locks on dependent segments. If no program holds locks on dependent segments, it is not necessary to lock dependent segments when they are accessed.

The following locking protocol allows IMS to make this determination. If a root segment is updated, the root lock is held at update level until commit. If a dependent segment is updated, it is locked at update level. When exiting the database record, the root segment is demoted to read level. When a program enters the database record and obtains the lock at either read or update level, the lock manager provides feedback indicating whether or not another program has the

lock at read level. This determines if dependent segments will be locked when they are accessed. For HISAM, the primary logical record is treated as the root, and the overflow logical records are treated as dependent segments.

Related Reading: For a special case involving the HISAM delete byte with parameter `ERASE=YES`, see “Deleting Segments” on page 76.

These lock protocols apply when the PI lock manager is used; however, if the IRLM is used, no lock is obtained when a dependent segment is updated. Instead, the root lock is held at single update level when exiting the database record. Therefore, no additional locks are required if a dependent segment is inserted, deleted, or replaced.

Locking for Q Command Codes

When a Q command code is issued for a root or dependent segment, a Q command code lock at share level is obtained for the segment. This lock is not released until a DEQ call with the same class is issued, or until commit time.

If a root segment is returned in hold status, the root lock obtained when entering the database record prevents another user with update capability from entering the database record. If a dependent segment is returned in hold status, a Q command code test lock is required. An indicator is turned on whenever a Q command code lock is issued for a database. This indicator is reset whenever the only application scheduled against the database ends. If the indicator is not set, then no Q command code locks are outstanding and no test lock is required to return a dependent segment in hold status.

Resource Locking Considerations with Block Level Sharing

Resource locking can occur either locally in a non-sysplex environment or globally in a sysplex environment.

In a non-sysplex environment, local locks can be granted in one of three ways:

- **Immediately** because of one of the following reasons:
 - IMS was able to get the required IRLM locks, and there is no other interest on this resource.
 - The request is compatible with other holders or waiters.
- **Asynchronously** because the request could not get the required IRLM latches and was suspended. (This can also occur in a sysplex environment.) The lock is granted when latches become available and one of three conditions exist:
 - No other holders exist.
 - The request is compatible with other holders or waiters.
 - The request is not compatible with the holders or waiters and was granted after their interest was released. (This could also occur in a sysplex environment.)

In a sysplex environment, global locks can be granted in one of three ways:

- **Locally by the IRLM** because of one of the following reasons:
 - There is no other interest for this resource.
 - This IRLM has the only interest, this request is compatible with the holders or waiters on this system, and XES already knows about the resource.
- **Synchronously on the XES CALL** because of one of the following reasons:
 - XES shows no other interest for this resource.
 - XES shows only SHARE interest for the hash class.

- **Asynchronously on the XES CALL** because of one of three conditions:
 - Either XES shows EXCLUSIVE interest on the hash class by an IRLM, but the resource names do not match (FALSE CONTENTION by RMF™).
 - Or XES shows EXCLUSIVE interest on the hash class by an IRLM and the resource names match, but the IRLM CONTENTION EXIT grants it anyway because the STATES are compatible (IRLM FALSE CONTENTION).
 - Or the request is incompatible with the other HOLDERS and is granted by the CONTENTION Exit after their interest is released (IRLM REAL CONTENTION).

Data Sharing Impact on Locking

When you use block-level data sharing, the IRLM must obtain the concurrence of the sharing system before granting global locks. Root locks are global locks, and dependent segment locks are not. When you use block-level data sharing, locks prevent the sharing systems from concurrently updating the same buffer. The buffer is locked before making the update, and the lock is held until after the buffer is written during commit processing. No buffer locks are obtained when a buffer is read.

If a Q command code is issued on any segment, the buffer is locked. This prevents the sharing system from updating the buffer until the Q command code lock is released.

Locking in HDAM, PHDAM, HIDAM, and PHIDAM Databases

If you access a HIDAM or PHIDAM root through the index, a lock is obtained on the index, using the RBA of the root segment as the resource name. Consequently, a single lock request locks both the index and the root.

When NOTWIN pointers are specified on a PHIDAM root, a lock on the next higher non-deleted root is required to provide data integrity. IMS obtains the additional lock by reading the index until a non-deleted index entry is found and then locking the RBA of the root segment as the resource name.

When you access an HDAM or a PHDAM database, the anchor of the desired root segment is locked as long as position exists on any root chained from that anchor. Therefore, if an update PCB has position on an HDAM or PHDAM root, no other user can access that anchor. When a segment has been updated and the IRLM is used, no other user can access the anchor until the user that is updating commits. If the PI lock manager is used and an uncommitted unit of work holds the anchor, locks are needed to access all root and dependent segments chained from the anchor until the user that is updating commits.

Locking for Secondary Indexes

When a secondary index is inserted, deleted or replaced, it is locked with a root segment lock. When the secondary index is used to access the target of the secondary index, depending on what the index points to, it might be necessary to lock the secondary index.

Managing I/O Errors

When a database I/O error occurs, IMS copies the buffer contents of the error block/control interval (CI) to a virtual buffer. A subsequent DL/I request causes the error block/CI to be read back into the buffer pool. The write error information and buffers are maintained across restarts, deferring recovery to a convenient time. I/O

Managing I/O Errors

error retry is automatically performed at database close time. If the retry is successful, the error condition no longer exists and recovery is not needed.

When a database I/O error occurs in a sysplex environment, the local system maintains the buffer and informs all members of the data-sharing group with registered interest in the database that the CI is unavailable. Subsequent DL/I requests for that CI receive a failure return code as long as the I/O error persists.

Although you do not have to register your databases with DBRC for error handling to work, registration is required for HALDBs and highly recommended for all other full-function databases.

If an error occurs on a database registered with DBRC and the system stops, the database could be damaged if the system is restarted and a /DBR command is not issued prior to accessing the database. The restart causes the error buffers to be restored as they were when the system stopped. If the same block had been updated during the batch run, the batch update would be overlaid.

Chapter 7. Choosing Fast Path Database Types

This chapter describes the characteristics and basic functions of Fast Path databases to help you decide what type of database to use. Fast Path databases include data entry databases (DEDBs) and main storage databases (MSDBs). DEDBs provide efficient storage for and access to large volumes of data. DEDBs also provide a high level of availability to that data. MSDBs store and provide access to an installation's most frequently used data.

Both DEDBs and MSDBs use the direct method of storing data. With the direct method, the hierarchic sequence of segments is maintained by putting direct-address pointers in each segment's prefix.

Each IMS environment supports Fast Path databases as follows:

- DB/DC supports both DEDBs and MSDBs.
- DBCTL supports DEDBs, but does not support MSDBs.
- DCCTL does not support DEDBs or MSDBs.

For a summary of the different characteristics of all IMS database types, including Fast Path databases, see Table 7 on page 63.

In this Chapter:

- "Data Entry Databases"
- "Main Storage Databases (MSDBs)" on page 132
- "Fast Path Virtual Storage Option" on page 139
- "Fast Path Synchronization Points" on page 153
- "Managing I/O Errors and Long Wait Times" on page 154

Data Entry Databases

Data entry databases (DEDBs) provide efficient storage for and access to large volumes of data. DEDBs also provide a high level of availability of that data.

Several characteristics of DEDBs also make DEDBs useful when you must gather detailed and summary information. These characteristics include:

- Area format
- Area data set replication
- Record deactivation
- Non-recovery option

A DEDB is a hierarchical database that contains up to 127 segment types: a root segment, an optional sequential dependent segment, and 0 to 126 direct dependent segments. If an optional sequential dependent segment is defined, you can define no more than 125 direct dependent segments. A DEDB structure can have as many as 15 hierarchical levels. Instances of sequential dependent segments for an area are stored in chronological order, regardless of the root on which they are dependent. Direct dependent segments are stored hierarchically, which allows for rapid retrieval.

Recommendation: Because ETO terminals cannot access terminal-related MSDBs, you should develop any new Fast Path databases as DEDBs instead of as MSDBs. You should also consider converting any of your existing non-terminal-related

Data Entry Databases

MSDBs with non-terminal-related keys to VSO DEDBs. You can use the MSDB-to-DEDB Conversion utility to do so.

Related Reading: For more information about the MSDB-to-DEDB Conversion utility (DBFUCDB0), see *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

DEDB Functions

DEDBs and MSDBs have many similar functions, including:

- Virtual storage
- The field (FLD) call
- Fixed length segments
- MSDB or DEDB commit view

In addition, DEDBs have the following functions and support:

- Full DBRC support
- Block-level sharing of areas available to
 - DBCTL
 - LU 6.2 applications
 - DB/DC applications
- RSR tracking
- HSSP support
- DEDB utilities
- Online database maintenance
- A full hierarchical model, including support of INSERT and DELETE calls
- A randomizer search technique

DEDB Areas

A DEDB can be organized into one or more data sets called areas. Areas increase the efficiency, capacity, and flexibility of DEDBs. This topic discusses DEDB areas and how to work with them.

Areas and the DEDB Format

The physical format of DEDBs makes the data they contain more readily available. In a hierarchical IMS database that does not use areas, the logical data structure is spread across the entire database. If multiple data sets are used, the data structure is broken up on a segment basis. A DEDB can use multiple data sets, called areas, with each area containing the entire data structure (see Figure 70 on page 128). Each area in a DEDB is a VSAM data set. A DEDB record (a root and its dependent segments) does not span areas. A DEDB can be divided into as many as 2048 such areas. This organization is transparent to the application program.

The randomizing module is used to determine which records are placed in each area. Because of the area concept, larger databases can exceed the limitation of 2^{32} bytes for a single VSAM data set. Each area can have its own space management parameters. You can choose these parameters according to the message volume, which can vary from area to area. DEDB areas can be allocated on different volume types.

Initialization, reorganization, and recovery of DEDBs are done on an area basis. Resource allocation is done at the control interval (CI) level. Multiple programs,

optionally together with one online utility, can access an area concurrently within a database, as long as they are using different CIs. CI sizes can be 512 bytes, 1K, 2K, 4K, and up to 28K in 4K increments. The media manager and Integrated Catalog Facility catalog of Data Facility Storage Management Subsystem (DFSMS) are required.

Opening and Preopening DEDB Areas

By default, IMS does not open a DEDB area until an eligible application accesses the area. Although this prevents unneeded areas from being opened at startup, it burdens the first application that accesses a DEDB area with some additional processing overhead. Multiple calls to multiple areas immediately following a startup process can increase this burden significantly.

You can limit the overhead of opening areas by preopening your DEDB areas. You can also distribute this overhead between the startup process and online operation by preopening only those areas that applications use the most and by leaving all other areas closed until an application first accesses them.

You specify the preopen status of an area using the PREOPEN and NOPREO parameters of the DBRC INIT.DBDS command or CHANGE.DBDS command.

By default IMS preopens all DEDB areas that have been assigned preopen status during the startup process; however, preopening a large number of DEDB areas during the startup process can delay data processing. To avoid this delay, you can have IMS preopen DEDB areas after the startup process and asynchronously to the execution of your application programs. In this case, if IMS has not preopened a DEDB area when an application program attempts to access the area, IMS opens the DEDB area at that time. You can specify this behavior by using the FPOPN= keyword in the IMS and DBC startup procedures. Specifically, FPOPN=P causes IMS to preopen DEDB areas after startup and asynchronous to application program execution.

The FPOPN= keyword determines how IMS reopens DEDB areas for both normal restarts (/NRE) and emergency restarts (/ERE).

Related Reading:

- For more information about the FPOPN= keyword and the IMS and DBC procedures, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.
- For more information about DBRC and DBRC commands, see the *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

Reopening DEDB Areas During an Emergency Restart: You can specify how IMS reopens DEDB areas during an emergency restart by using the FPOPN= keyword in the IMS procedure or DBC procedure. The following list describes how the FPOPN= keyword affects the reopening of DEDB areas during an emergency restart:

FPOPN=N

During the startup process, IMS opens only those areas that have preopen status. This is the default.

FPOPN=P

After the startup process completes and asynchronous to the resumption of application processing, IMS opens only those areas that have preopen status.

FPOPN=R

After the startup process completes and asynchronous to the resumption of application processing, IMS opens only those areas that were open prior to the abnormal termination. All DEDB areas that were closed at the time of the abnormal termination, including DEDB areas with a preopen status, will remain closed when you restart IMS.

FPOPN=D

Suppresses the preopen process. DEDB areas that have a preopen status are not preopened and remain closed until they are first accessed by an application program or until they are manually opened with a `/START AREA` command.

FPOPN=D overrides, but does not change, the preopen status of DEDB areas as set by the PREOPEN parameter of the DBRC commands `INIT.DBDS` and `CHANGE.DBDS`.

Related Reading: For more information about the FPOPN= keyword and the IMS and DBC startup procedures, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Stopping and Starting DEDBs and DEDB Areas

You can prevent access to a DEDB by stopping it with the `/STOP DATABASE` command. You can also prevent access to a single DEDB area by stopping it with the `/STOP AREA` command. These commands do not affect application programs currently scheduled against the DEDB, but prevent IMS from scheduling any new application programs that need access to the stopped database or area.

You can resume access to a stopped DEDB by starting it with the `/START DATABASE` command. You can also resume access to a stopped area by starting it with the `/START AREA` command. The `/START AREA` command does not open areas unless you have specified them as PREOPEN areas.

Restarting and Reopening Areas After an IRLM Failure

The internal resource lock manager (IRLM) ensures the integrity of databases in a data sharing environment. To avoid compromising the integrity of the data in DEDB areas when an IRLM fails, all DEDB areas under the control of the failed IRLM are stopped. After you correct the failure and reconnect IRLM to the IMS system, you must restart and reopen the DEDB areas that the IRLM controls.

You can specify how IMS restarts and reopens DEDB areas after the IRLM reconnects, by using the FPRLM= keyword in the IMS and DBC procedures. The following list describes how the FPRLM= keyword affects the reopening of DEDB areas after an IRLM failure has been corrected:

FPRLM=N

All DEDB areas remain stopped and unopened until you issue a `/START DATABASE` or `/START AREA` command. This is the default.

FPRLM=S

After IRLM reconnects, IMS restarts, but does not reopen, all areas that were open at the time of the IRLM failure. IMS restarts the DEDB areas asynchronously to the resumption of application processing.

FPRLM=R

After IRLM reconnects, IMS restores all DEDB areas to their state at the time of the IRLM failure, restarting and reopening DEDB areas regardless of whether the DEDB areas have preopen status. IMS restores the DEDB areas asynchronously to the resumption of application processing.

FPRLM=A

After IRLM reconnects, IMS restarts and reopens all DEDB areas that were open at the time of the IRLM failure and all DEDB areas that have a preopen status, even if they were closed at the time of the IRLM failure. IMS restores the DEDB areas asynchronously to the resumption of application processing.

Related Reading:

- For more information about the FPRLM= keyword and the IMS and DBC procedures, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.
- For more information about IRLM, see:
 - *IMS Version 9: Operations Guide*
 - *IMS Version 9: Administration Guide: System*

Read and Write Errors in DEDB Areas

This topic describes how IMS handles read and write errors that occur in DEDB areas.

Read Error: When a read error is detected in an area, the application program receives an AO status code. An Error Queue Element (EQE) is created, but not written to the second CI nor sent to the sharing system in a data sharing environment. Application programs can continue to access that area; they are prevented only from accessing the CI in error. After read errors on four different CIs, the area data set (ADS) is stopped. The read errors must be consecutive; that is, if there is an intervening write error, the read EQE count is cleared. This read error processing only applies to a multiple area data set (MADS) environment.

Write Error: When a write error is detected in an area, an EQE is created and application programs are allowed access to the area until the EQE count reaches 11. Even though part of a database might not be available (one or more areas are stopped), the database is still logically available and transactions using that database are still scheduled. If multiple data sets make up the area, chances are that one copy of the data will always be available.

If your DEDB is nonrecoverable, write errors are handled differently, compared to recoverable DEDBs. When there is a write error in an area, an EQE is created. When there are 10 EQEs for an area, DBRC marks it "Recovery Needed" and IMS stops the area. If the area is shared, then all IMSs in the sharing group are notified and they also stop the area. When a DEDB is marked "Recovery Needed", you must restore it, such as from an image copy. Incorporate this recovery procedure into your operational procedures.

When a write error occurs to a DEDB using MADS, an EQE is created for the ADS that had the write error. In this environment, when the maximum of 10 EQEs is reached, the ADS is stopped.

When a write error to a recoverable DEDB area using a single ADS occurs, IMS invokes the I/O toleration (IOT) processing. IMS allocates a virtual buffer in ECSA and copies the control interval in error from the Fast Path common buffer to the virtual buffer. IMS records the creation of the virtual buffer with an X'26' log record. If the database is registered with DBRC, an Extended Error Queue Element (EEQE) is created and registered in DBRC. The EEQE identifies the control interval in error. In a data sharing environment using IRLM, all sharing partners are notified of the creation of the EEQE.

The data that is tolerated is available to the IMS system that created the EEQE. The sharing partner will get an 'AO' status when it requests that CI because the data is not available. When a request is made for a control interval that is tolerated, the data is copied from the virtual buffer to a common buffer. When an update is performed on the data, it is copied back to the virtual buffer. A standard X'5950' log record is generated for the update.

Every write error is represented by an EEQE on an area basis. The EEQEs are maintained by DBRC and logged to the IMS log as X'26' log records. There is no logical limit to the number of EEQEs that can exist for an area. There is a physical storage limitation in DBRC and ECSA for the number of EEQEs that can be maintained. This limit is installation dependent. To make sure that we do not overextend DBRC or ECSA usage, a limited number of EEQEs are allowed for a DEDB. The limit is 100. After 100 EEQEs are created for an area, the area is stopped.

During system checkpoint, /ST0, and /VUN commands, IMS attempts to write back the CIs in error. If the write is successful, the EEQE is removed. If the write is unsuccessful, the EEQE remains.

Record Deactivation

If an error occurs while an application program is updating a DEDB, it is not necessary to stop the database or even the area. IMS continues to allow application programs to access that area. It only prevents them from accessing the control interval in error by creating an EQE for the error CI. If there are multiple copies of the area, chances are that one copy of the data will always be available. It is unlikely that the same control interval will be in error in all copies of the area. IMS automatically makes an area data set unavailable when a count of 11 errors has been reached for that data set.

Record deactivation minimizes the effect of database failure and errors to the data in these ways:

- If multiple copies of an area data set are used, and an error occurs while an application program is trying to update that area, the error does not need to be corrected immediately. Other application programs can continue to access the data in that area through other available copies of that area.
- If a copy of an area has a number of I/O errors, you can create a new copy from existing copies of the area using the DEDB Area Data Set Create utility. The copy with the errors can then be destroyed.

Non-Recovery Option

By specifying a VSO or non-VSO DEDB as nonrecoverable, you can improve online performance and reduce database change logging of your DEDBs. IMS does not log any changes from a nonrecoverable DEDB, nor does it keep any updates in the DBRC RECON data set. All areas are nonrecoverable in a nonrecoverable DEDB.

SDEPs are not supported for nonrecoverable DEDBs. After IMS calls DBRC to authorize the areas, IMS checks for SDEPs. If IMS finds SDEPs, IMS calls DBRC to unauthorize them and IMS stops them. You must remove the SDEP segment type from the DEDB design before IMS will authorize the DEDB.

Unlike full-function nonrecoverable databases, which support backout, nonrecoverable DEDBs are truly nonrecoverable and cannot REDO during restart or XRF takeover. IMS writes a single log record, x'5951', once for every area at each sync point to indicate that nonrecoverable suppression has taken place. The x'5951' log and DMAC flags determine the integrity of an area during an emergency restart

or XRF takeover. If there are errors found in a nonrecoverable DEDB during an XRF takeover or an emergency restart, message DFS3711W is issued and the DEDB is not stopped.

Related Reading: For information on how IMS handles nonrecoverable DEDB write errors, which can happen during restart or XRF takeover, see “Write Error” on page 117.

Nonrecoverable DEDBs must register with DBRC. To define a DEDB as nonrecoverable, use the command `INIT.DB DBD() TYPEFP NONRECOV`. The default is `RECOVABL` for recoverable DEDB.

Before changing the recoverability of a DEDB, issue a `/STOP DB`, `/STO AREA`, or `/DBR DB` command. To change a recoverable DEDB to a nonrecoverable DEDB, use the DBRC command `CHANGE.DB DBD() NONRECOV`. To change nonrecoverable DEDB to a recoverable DEDB, use the command `CHANGE.DB DBD() RECOVABL`.

To restore a nonrecoverable DEDB, use the `GENJCL.RECOV RESTORE` command. The recovery utility restores the database to the last image copy taken. If the DEDB had been changed from a recoverable DEDB to a nonrecoverable DEDB, the recovery utility will apply any updates from the logs up to the point when the change was made (if no image copy was made after the change to nonrecoverable).

Area Data Set Replication

A data set can be copied, or replicated, up to seven times, increasing the availability of the data to application programs. The DEDB Area Data Set Create utility (`DBFUMRI0`) produces one or more copies of a data set representing the area without stopping the area. All copies of an area data set must have identical CI sizes and spaces but can reside on different devices. The utility uses all the current copies to complete its new data set, proceeding to another copy if it detects an I/O error for a particular record. In this way, a clean copy is constructed from the aggregate of the available data. Current updates to the new data set take effect immediately.

The Create utility can create its new copy on a different device, as specified in its job control language (JCL). If your installation was migrating data to other storage devices, then this process could be carried out while the online system was still executing, and the data would remain current.

To ensure all copies of a DEDB remain identical, IMS updates all copies when a change is made to only one copy.

If an ADS fails open during normal open processing of a DEDB with multiple data sets (MADS), none of the copies of the ADS can be allocated, and the area is stopped. However, when open failure occurs during emergency restart, only the failed ADS is unallocated and stopped. The other copies of the ADS remain available for use.

DEDBs and Data Sharing

You can specify different levels of data sharing for DEDBs. The specifications you make for a DEDB apply to all of the areas the DEDB contains.

If you specify that a DEDB does not allow data sharing, only one IMS system can access a DEDB area at a time; however, other IMS systems can still access the other areas the DEDB contains.

Data Entry Databases

If you specify that a DEDB allows data sharing, multiple IMS systems can access the same DEDB area at the same time. Sharing a single DEDB area is equivalent to block-level sharing of full-function databases.

You can specify the level of data sharing that a DEDB allows by using the SHARELVL parameter in the DBRC commands INIT.DB and CHANGE.DB. If any IMS has already authorized the database, changing the SHARELVL does not modify the database record. The SHARELVL parameter applies to all areas in a DEDB.

You can share DEDB areas directly from DASD or from a coupling facility structure using the Virtual Storage Option (VSO).

Related Reading:

- For general information on VSO, including its benefits and use, see “Fast Path Virtual Storage Option” on page 139.
- For specific information on sharing VSO DEDB areas, see “Sharing of VSO DEDB Areas” on page 142.
- For more information on the SHARELVL parameter, see the *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.
- For general information on data sharing, see *IMS Version 9: Administration Guide: System*.

Fixed- and Variable-Length Segments in DEDBs

DEDBs support fixed-length segments. Thus you can define fixed-length or variable-length segments for your DEDBs. This support allows you to use MSDB applications for your DEDBs.

To define fixed-length segments, specify a single value for the BYTES= parameter during DBDGEN in the SEGM macro. To define variable-length segments, specify two values for the BYTES= parameter during DBDGEN in the SEGM macro.

Application programs for fixed-length-segment DEDBs, like MSDBs, do not see the length (LL) field at the beginning of each segment. Application programs for variable-length-segment DEDBs do see the length (LL) field at the beginning of each segment, and must use it to process the segment properly.

Fixed-length-segment application programs using REPL and ISRT calls can omit the length (LL) field.

Examples of Defining Segments

Figure 63 and Figure 64 show examples of how to use the BYTES= parameter to define variable-length or fixed-length segments.

```
ROOTSEG  SEGM  NAME=ROOTSEG1,          C
              PARENT=0,                  C
              BYTES=(390,20)
```

Figure 63. Defining a Variable-Length Segment

```
ROOTSEG  SEGM  NAME=ROOTSEG1,          C
              PARENT=0,                  C
              BYTES=(320)
```

Figure 64. Defining a Fixed-Length Segment

Parts of a DEDB Area

A DEDB area consists of three parts:

- Root addressable part
- Independent overflow part
- Sequential dependent part

Figure 65 on page 122 shows these parts of a DEDB area. Each part is described in detail in the following topics:

- “Root Addressable Part” on page 123
- “Independent Overflow Part” on page 123
- “Sequential Dependent Part” on page 123
- “CI and Segment Formats” on page 123

Data Entry Databases




Figure 65. Parts of a DEDB Area in Storage

When a DEDB data set is initialized by the DEDB initialization utility (DBFUMIN0), additional CIs are created for internal use, so the DEDB area will actually contain more CIs than are shown in Figure 65. These extra CIs are referred to as the Reorganization UOW and were used for the DEDB Direct Reorganization utility (DBFUMDR0), which went out service with IMS Version 5 and was replaced by the High-Speed DEDB Direct Reorganization utility (DBFUHDR0). Although IMS does not use the extra CIs, DBFUMIN0 creates them for compatibility purposes.

Root Addressable Part

The root addressable part is divided into units-of-work (UOW), which are the basic elements of space allocation. A UOW consists of a user-specified number of CIs located physically contiguous.

Each UOW in the root addressable part is further divided into a base section and an overflow section. The base section contains CIs of a UOW that are addressed by the randomizing module, whereas the overflow section of the UOW is used as a logical extension of a CI within that UOW.

Root and direct dependent segments are stored in the base section. Both can be stored in the overflow section if the base section is full.

Independent Overflow Part

The independent overflow part contains empty CIs that can be used by any UOW in the area. When a UOW gets a CI from the independent overflow part, the CI can be used only by that UOW. A CI in the independent overflow part can be considered an extension of the overflow section in the root addressable part as soon as it is allocated to a UOW. The independent overflow CI remains allocated to a specific UOW unless, after a reorganization, it is no longer required, at which time it is freed.

Sequential Dependent Part

The sequential dependent part holds sequential dependent segments from roots in all UOWs in the area. Sequential dependent segments are stored in chronological order without regard to the root or UOW that contains the root. When the sequential dependent part is full, it is reused from the beginning. However, before the sequential dependent part can be reused, you must use the Delete utility DBFUMDLO to delete a contiguous portion, or all the sequential dependent segments in that part.

CI and Segment Formats

This topic contains diagnosis, modification, or tuning information.

The following four diagrams—Figure 66, Figure 67 on page 124, Figure 68 on page 125, and Figure 69 on page 125—show the following formats:

- CI format
- Root segment format
- Sequential dependent segment format
- Direct dependent segment format

The tables that follow each figure—Table 10 on page 124, Table 11 on page 124, Table 12 on page 125, and Table 13 on page 125, respectively—describe the sections of the CI and segments in the order that the sections appear in the graphic.




Figure 66. CI Format

Table 10. CI Format

| CI Section | Number of Bytes | Explanation |
|------------|-----------------|---|
| FSE AP | 2 bytes | Offset to the first free space element. These 2 bytes are unused if the CI is in the sequential dependent part. |
| CI TYP | 2 bytes | Describes the use of this CI and the meaning of the next 4 bytes. |
| RAP | 4 bytes | Root anchor point if this CI belongs to the base section of the root addressable area. All root segments randomizing to this CI are chained off this RAP in ascending key sequence. Only one RAP exists per CI. Attention: In the dependent and independent overflow parts, these 4 bytes are used by Fast Path control information. No RAP exists in sequential dependent CIs. |
| CUSN | 2 bytes | CI Update Sequence Number (CUSN). A sequence number maintained in each CI. It is increased with each update of the particular CI during the synchronization process. |
| RBA | 4 bytes | Relative byte address of this CI. |
| RDF | 3 bytes | Record definition field (contains VSAM control information). |
| CIDF | 4 bytes | CI definition field (contains VSAM control information). |




Figure 67. Root Segment Format (with Sequential and Direct Dependent Segments with Subset Pointers)

Table 11. Root Segment Format

| Segment Section | Number of Bytes | Explanation |
|-----------------|-----------------|--|
| SC | 1 byte | Segment code. |
| PD | 1 byte | Prefix descriptor. |
| PTF | 4 bytes | Physical twin forward pointer. Contains the RBA of the next root in key sequence. |
| SPCF | 8 bytes | Sequential physical child first pointer. Contains the cycle count and RBA of the last inserted sequential dependent under this root. This pointer will not exist if the sequential dependent segment is not defined. |
| PCF | 4 bytes | Physical child first pointer. PCF points to the first occurrence of a direct dependent segment type. There can be up to 126 PCF pointers or 125 PCF pointers if there is a sequential dependent segment. PCF pointers will not exist if direct dependent segments are not defined. |
| PCL | 4 bytes | Physical child last pointer. PCL is an optional pointer that points to the last physical child of a segment type. This pointer will not exist if direct dependent segments are not defined. |

Table 11. Root Segment Format (continued)

| Segment Section | Number of Bytes | Explanation |
|-----------------|-----------------|--|
| SSP | 4 bytes | Subset pointer. For each child type of the parent, up to eight optional subset pointers can exist. |
| LL | 2 bytes | Variable length of this segment. |




Figure 68. Sequential Dependent Segment Format

Table 12. Sequential Dependent Segment Format

| Segment Section | Number of Bytes | Explanation |
|-----------------|-----------------|---|
| SC | 1 byte | Segment code. |
| UN | 1 byte | Prefix descriptor. |
| SPTF | 8 bytes | Sequential physical twin forward pointer. Contains the cycle count and the RBA of the immediately preceding sequential dependent segment under the same root. |
| LL | 2 bytes | Variable length of this segment. |




Figure 69. Direct Dependent Segment Format

Table 13. Direct Dependent Segment Format

| Segment Section | Number of Bytes | Explanation |
|-----------------|-----------------|--|
| SC | 1 byte | Segment code. |
| UN | 1 byte | Unused. |
| PTF | 4 bytes | Physical twin forward pointer. Contains the RBA of the next occurrence of this direct dependent segment type. |
| PCF | 4 bytes | Physical child first pointer. PCF points to the first occurrence of a direct dependent segment type. In a direct dependent segment there can be up to 125 PCF pointers or 124 PCF pointers if there is a sequential dependent segment. PCF pointers will not exist if direct dependent segments are not defined. |

Table 13. Direct Dependent Segment Format (continued)

| Segment Section | Number of Bytes | Explanation |
|-----------------|-----------------|---|
| PCL | 4 bytes | Physical child last pointer. PCL is an optional pointer that points to the last physical child of a segment type. This pointer will not exist if direct dependent segments are not defined. |
| SSP | 4 bytes | Subset pointer. For each child type of the parent, up to eight optional subset pointers can exist. |
| LL | 2 bytes | Variable length of this segment. |

Root Segment Storage

DEDB root segments are stored as prescribed by the randomizing routine, and are chained in ascending key sequence from each anchor point.

Related Reading: For information on the system-supplied or user-supplied randomizing module for DEDBs, see *IMS Version 9: Customization Guide*.

Each CI in the base section of a UOW in an area has a single anchor point. Sequential processing using GN calls processes the roots in the following order:

1. Ascending area number
2. Ascending UOW
3. Ascending key in each anchor point chain

Each root segment contains, in ascending key sequence, a PTF pointer containing the RBA of the next root.

Direct Dependent Segment Storage

The DEDB maintains processing efficiency while supporting a hierarchic physical structure with direct dependent segment types. A maximum of 127 segment types are supported (up to 126 direct dependent segment types, or 125 if a sequential dependent segment is present).

Direct dependent (DDEP) segment types can be efficiently retrieved hierarchically, and the user has complete online processing control over the segments. Supported processing options are insert, get, delete, and replace. With the replace function, users can alter the length of the segment. DEDB space management logic attempts to store an inserted direct dependent in the same CI that contains its root segment. If insufficient space is available in that CI, the root addressable overflow and then the independent overflow portion of the area are searched.

DDEP segments can be defined with or without a unique sequence field, and are stored in ascending key sequence.

Physical chaining of direct dependent segments consists of a physical child first (PCF) pointer in the parent for each defined dependent segment type and a physical twin forward (PTF) pointer in each dependent segment.

DEDBs allow a PCL pointer to be used. This pointer makes it possible to access the last physical child of a segment type directly from the physical parent. The INSERT rule LAST avoids the need to follow a potentially long physical child pointer chain.

Subset pointers are a means of dividing a chain of segment occurrences under the same parent into two or more groups, of subsets. You can define as many as eight subset pointers for any segment type, dividing the chain into as many as nine subsets. Each subset pointer points to the start of a new subset.

Related Reading: For more information on defining and using subset pointers, see *IMS Version 9: Application Programming: Database Manager*.

Sequential Dependent Segment Storage

DEDB sequential dependent (SDEP) segments are stored in the sequential dependent part of an area in the order of entry. SDEP segments chained from different roots in an area are intermixed in the sequential part of an area without regard to which roots are their parents. SDEP segments are specifically designed for fast insert capability. However, online retrieval is not as efficient because increased input operations can result.

If all SDEP dependents were chained from a single root segment, processing with Get Next in Parent calls would result in a backward sequential order. (Some applications are able to use this method.) Normally, SDEP segments are retrieved sequentially only by using the DEDB Sequential Dependent Scan utility (DBFUMSCO), described in *IMS Version 9: Utilities Reference: Database and Transaction Manager*. SDEP segments are then processed by offline jobs.

SDEP segments are used for data collection, journaling, and auditing applications.

Enqueue Level of Segment CIs

Allocation of CIs involves three different enqueue levels.

- A NO ENQ level, which is typical of any SDEP CI. SDEP segments can never be updated; therefore they can be accessed and shared by all regions at the same time.
- A SHARED level, which means that the CI can be shared between non-update transactions. A CI at the SHARED level delays requests from any update transactions.
- An EXCLUSIVE level, which prevents contenders from acquiring the same resource.

The level of enqueue at which ROOT and SDEP segment CIs are originally acquired depends on the intent of the transaction. If the intent is update, all acquired CIs (with the exception of SDEP CIs) are held at the EXCLUSIVE level. If the intent is not update, they're held at the SHARED level. Of course, there is the potential for deadlock.

The level of enqueue, just described, is reexamined each time the buffer stealing facility is invoked. The buffer stealing facility examines each buffer (and CI) that is already allocated and updates its level of enqueue.

All other enqueued CIs are released and therefore can be allocated by other regions.

Related Reading: For more information about the buffer stealing facility, see "Fast Path Buffer Allocation Algorithm" on page 287.

Figure 70 on page 128 shows an example of DEDB structure.

Data Entry Databases

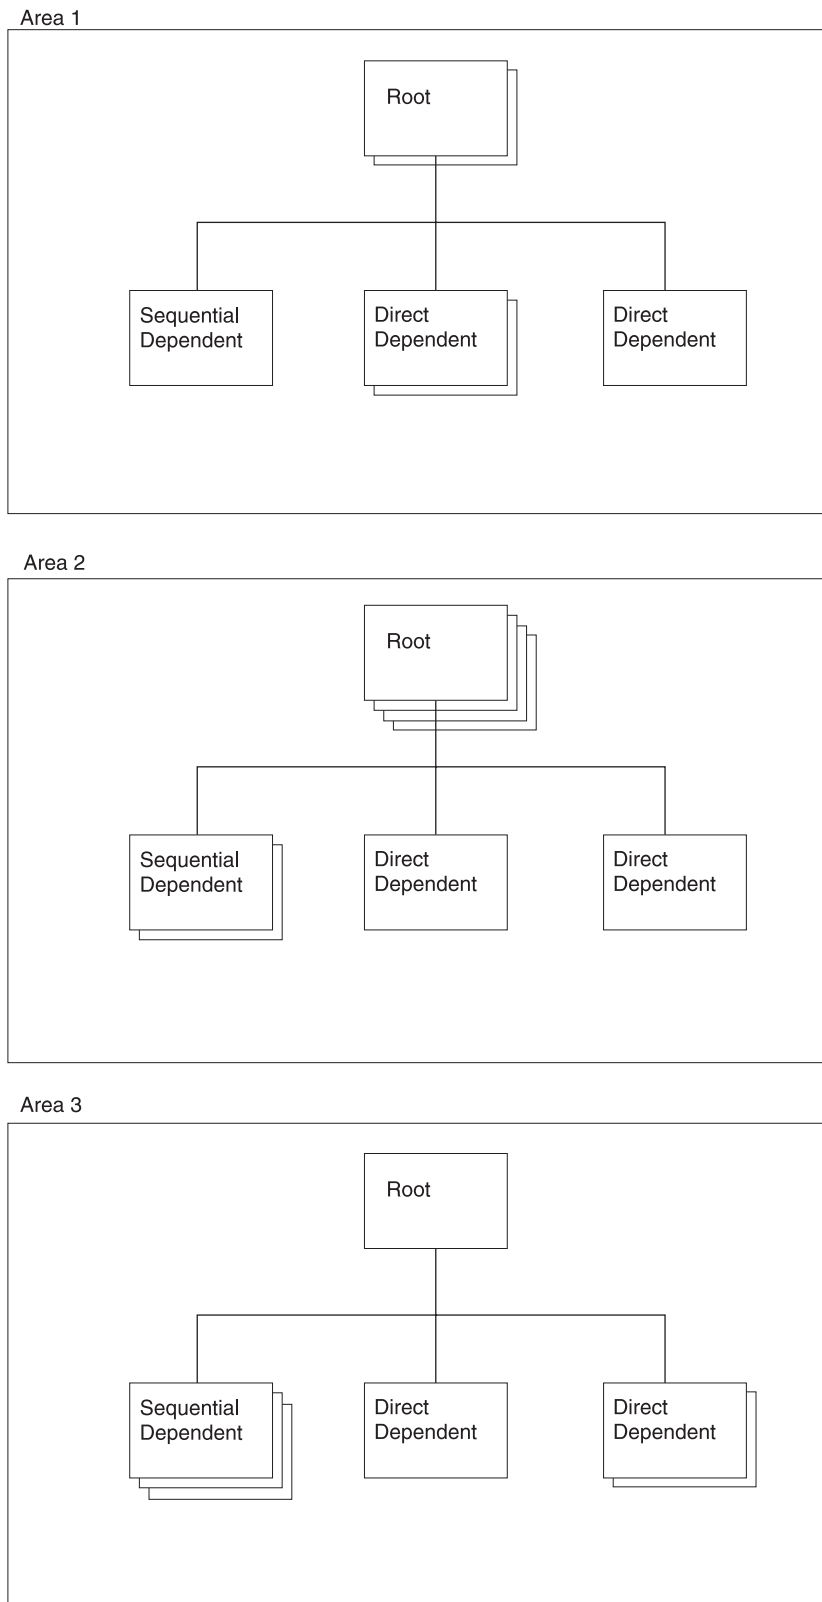


Figure 70. DEDB Structure Example

DEDB Space Search Algorithm

This topic contains diagnosis, modification, or tuning information.

The general rule for inserting a segment into a DEDB is the same as it is for an HD database. The rule is to store the segment (root and direct dependents) into the most desirable block.

For root segments, the most desirable block is the RAP CI. For direct dependents, the most desirable block is the root CI. When space for storing either roots or direct dependents is not available in the most desirable block, the DEDB insert algorithm (described next) searches for additional space. Space to store a segment could exist:

- In the dependent overflow
- In an independent overflow CI currently owned by this UOW

Additional independent overflow CIs would be allocated if required.

This algorithm attempts to store the data in the minimum amount of CIs rather than scatter database record segments across a greater number of RAP and overflow CIs. The trade-off is improved performance for future database record access versus optimum space utilization.

DEDB Insert Algorithm

This topic contains diagnosis, modification or tuning information.

The DEDB insert algorithm searches for additional space when space is not available in the most desirable block. For root segments, if the RAP CI does not have sufficient space to hold the entire record, it contains the root and as many direct dependents as possible. Base CIs that are not randomizer targets go unused. The algorithm next searches for space in the first dependent overflow CI for this UOW. From the header of the first dependent overflow CI, a determination is made whether space exists in that CI.

Related Reading: For information on DEDB CI format and allocation, see *IMS Version 9: Diagnosis Guide and Reference*.

If the CI pointed to by the current overflow pointer does not have enough space, the next dependent overflow CI (if one exists) is searched for space. The current overflow pointer is updated to point to this dependent overflow CI. If no more dependent overflow CIs are available, then the algorithm searches for space in the independent overflow part.

When an independent overflow CI has been selected for storing data, it can be considered a logical extension of the overflow part for the UOW that requested it.

Figure 71 on page 130 shows how a UOW is extended into independent overflow. This UOW, defined as 10 CIs, includes 8 Base CIs and 2 dependent overflow CIs. Additional space is needed to store the database records that randomize to this UOW. Two independent overflow CIs have been acquired, extending the size of this UOW to 12 CIs. The first dependent overflow CI has a pointer to the second independent overflow CI indicating that CI is the next place to look for space.




Figure 71. Extending a UOW to Use Independent Overflow

DEDB Free Space Algorithm

This topic contains diagnosis, modification, or tuning information.

The DEDB free space algorithm is used to free dependent overflow and independent overflow CIs. When a dependent overflow CI becomes entirely empty, it becomes the CI pointed to by the current overflow pointer in the first dependent overflow CI, indicating that this is the first overflow CI to use for overflow space if the most desirable block is full. An independent overflow CI is owned by the UOW to which it was allocated until every segment stored in it has been removed. When the last segment in an independent overflow CI is deleted, the empty CI is made available for reuse. When the last segment in a dependent overflow CI is deleted, it can be reused as described at the beginning of this topic.

A dependent overflow or an independent overflow CI can be freed in two ways:

- Reorganization
- Segment deletion

Reorganization

During online reorganization, the segments within a UOW are read in GN order and written to the reorganization utility private buffer set. This process inserts segments into the reorganization utility private buffer set, eliminating embedded free space. If all the segments do not fit into the reorganization utility private buffer set (RAP CI plus dependent overflow CIs), then new independent overflow CIs are allocated as needed. When the data in the reorganization utility private buffer set is copied back to the correct location, then the following events occur:

- The newly acquired independent overflow CIs are retained.
- The old segments are deleted.
- Previously allocated independent overflow CIs are freed.

Segment Deletion

A segment is deleted either by an application DLET call or because a segment is REPLaced with a different length. Segment REPLace can cause a segment to move. Full Function handles segment length increases differently from DEDBs. In Full Function, an increased segment length that does not fit into the available free space is split, and the data is inserted away from the prefix. For DEDBs, if the replaced segment is changed, it is first deleted and then reinserted. The insertion process follows the normal space allocation rules.

The REPL call can cause a dependent overflow or an independent overflow CI to be freed if the last segment is deleted from the CI.

Managing Unusable Space with IMS Tools

Space in a DEDB should be closely monitored to avoid out-of-space conditions for an area. Products such as the IMS High Performance (HP) Pointer Checker, which includes the Hierarchical Database (HD) Tuning Aid and Space Monitor tools, can identify the different percentages of free space in the RAP, dependent overflow, and independent overflow CIs. If a large amount of space exists in the RAP CIs or dependent overflow CIs, and independent overflow has a high use percentage, a reorganization can allow the data to be stored in the root addressable part, freeing up independent overflow CIs for use by other UOWs. The IMS HP Pointer Checker and the tools it includes can help you determine if the data distribution is reasonable.

For more information on tuning DEDBs, see “Tuning Fast Path Systems” on page 421.

DL/I Calls against a DEDB

This topic contains diagnosis, modification, or tuning information.

DEDB processing uses the same call interface as DL/I processing. Therefore, any DL/I call or calling sequence executed against a DEDB has the same logical result as if executed against an HDAM or PHDAM database.

The SSA rules for DEDBs have the following restrictions:

- You cannot use the Q command code with DEDBs.
- IMS ignores command codes used with sequential dependent segments.
- If you use the D command code in a call to a DEDB, the P processing option need not be specified in the PCB for the program. The P processing option has a different meaning for DEDBs than for DL/I databases.

Related Reading: For more information on how DEDBs are processed, see *IMS Version 9: Application Programming: Database Manager*.

Mixed Mode Processing

IMS application programs can run as message processing programs (MPPs), batch message processing programs (BMPs), and Fast Path programs (IFPs). IFPs can access full function databases. Similarly, MPPs and BMPs can access DEDBs and MSDBs.

Because of differences in sync point processing, there are differences in the way database updates are committed. IFPs that request full function resources, or MPPs

(or BMPs) that request DEDB (or MSDB) resources operate in “mixed mode”. The performance and resource use implications are discussed in “Fast Path Synchronization Points” on page 153.

Main Storage Databases (MSDBs)

The MSDB structure consists of fixed-length root segments only, although the root segment length can vary between MSDBs. The maximum length of any segment is 32000 bytes with a maximum key length of 240 bytes. Additional prefix data extends the maximum total record size to 32258 bytes.

The following options are *not* available for MSDBs:

- Multiple data set groups
- Logical relationships
- Secondary indexing
- Variable-length segments
- Field-level sensitivity

The MSDB family of databases consists of three types:

- Terminal-related *fixed* database
- Terminal-related *dynamic* database
- Non-terminal-related database *without* terminal keys

Recommendation: ETO terminals cannot access terminal-related MSDBs. IBM recommends that any new Fast Path database that you develop be DEDBs instead of MSDBs. Also, you should consider converting any of your existing non-terminal-related MSDBs with non-terminal-related-keys to VSO DEDBs. You can use the MSDB-to-DEDB Conversion utility.

An MSDB is defined in the DBD in the same way as any other IMS database, by coding ACCESS=MSDB in the DBD statement. The REL keyword in the DATASET statement selects one of the four MSDB types.

Both dynamic and fixed terminal-related MSDBs have the following characteristics:

- The record can be updated only through processing of messages issued from the LTERM that owns the record. However, the record can be read using messages from any LTERM.
- The name of the LTERM that owns a segment is the key of the segment. An LTERM cannot own more than one segment in any one MSDB.
- The key does not reside in the stored segment.
- Each segment in a fixed terminal-related MSDB is assigned to and owned by a different LTERM.

Non-terminal-related MSDBs have the following characteristics:

- No ownership of segments exists.
- No insert or delete calls are allowed.
- The key of segments can be an LTERM name or a field in the segment. As with a terminal-related MSDB, if the key is an LTERM name, it does not reside in the segment. If the key is not an LTERM name, it resides in the sequence field of the segment. If the key resides in the segment, the segments must be loaded in key sequence because, when a qualified SSA is issued on the key field, a binary search is initiated.

When to Use an MSDB

MSDBs store and provide access to an installation's most frequently used data. The data in an MSDB is stored in segments, and each segment available to one or all terminals.

MSDBs provide a high degree of parallelism and are suitable for applications in the banking industry (such as general ledger). To provide fast access and allow frequent update to this data, MSDBs reside in virtual storage during execution.

One use for a terminal-related fixed MSDB is in an application in which each segment contains data associated with a logical terminal. In this type of application, the application program can read the data (possibly for general reporting purposes) but cannot update it.

Non-terminal-related MSDBs (without terminal-related keys) are typically used in applications in which a large number of people need to update data at a high transaction rate. An example of this is a real-time inventory control application, in which reduction of inventory is noted from many cash registers.

MSDBs Storage

The MSDB Maintenance utility (DBFDBMA0) creates the MSDBINIT sequential data set in physical ascending sequence (see Figure 73 on page 134). During a cold start, or by operator request during a normal warm start, the sequential data set MSDBINIT is read and the MSDBs are created in virtual storage (see Figure 72).

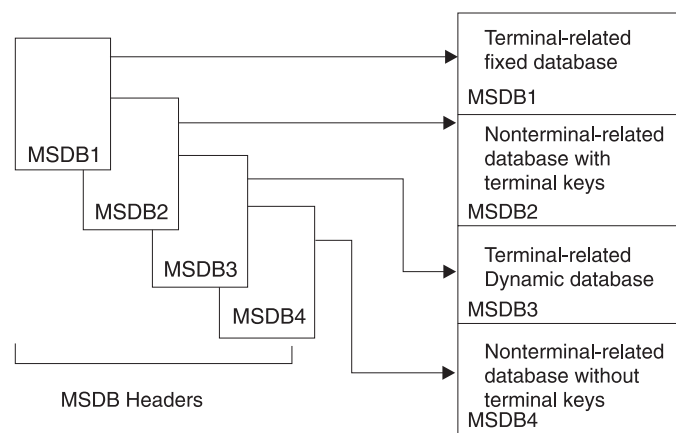


Figure 72. MSDB Pointers

During a warm start, the control program uses the current checkpoint data set for initialization. The MSDB Maintenance utility can also modify the contents of an old MSDBINIT data set. For warm start, the master terminal operator can request use of the IMS.MSDBINIT, rather than a checkpoint data set.

Diagnosis, Modification or Tuning Information

Figure 73 shows the MSDBINIT record format. Table 14 on page 134 explains the record parts.

Main Storage Databases (MSDBs)

| | LL | 00 | DBDname | Count | Type | KL | KEY | MSDB Segment |
|-------|----|----|---------|-------|------|----|--------|--------------------|
| Bytes | 2 | 2 | 8 | 4 | 1 | 1 | Varies | Varies(MAX 32,000) |

Figure 73. MSDBINIT Record Format

Table 14. MSDBINIT Record Format

| Record Part | Bytes | Explanation |
|---|--------|---|
| LL | 2 | Record length (32,258 maximum) |
| X'00' | 2 | Always hexadecimal zeros |
| DBDname | 8 | DBD name |
| Count | 4 | Segment count |
| Type | 1 | MSDB type: <ul style="list-style-type: none">• X'11' non-related• X'31' non-related with terminal keys• X'33' fixed related• X'37' dynamic related |
| KL | 1 | Key length (240 maximum) |
| Key | varies | Key or terminal name |
| MSDB segment | varies | MSDB segment (32,000 maximum) |
| _____ End of Diagnosis, Modification or Tuning Information _____ | | |

MSDB Record Storage

This topic contains diagnosis, modification, or tuning information.

MSDB records contain no pointers except the forward chain pointer (FCP) connecting free segment records in the terminal-related dynamic database.

Figure 74 on page 135 shows a high-level view of how MSDBs are arranged in priority sequence.

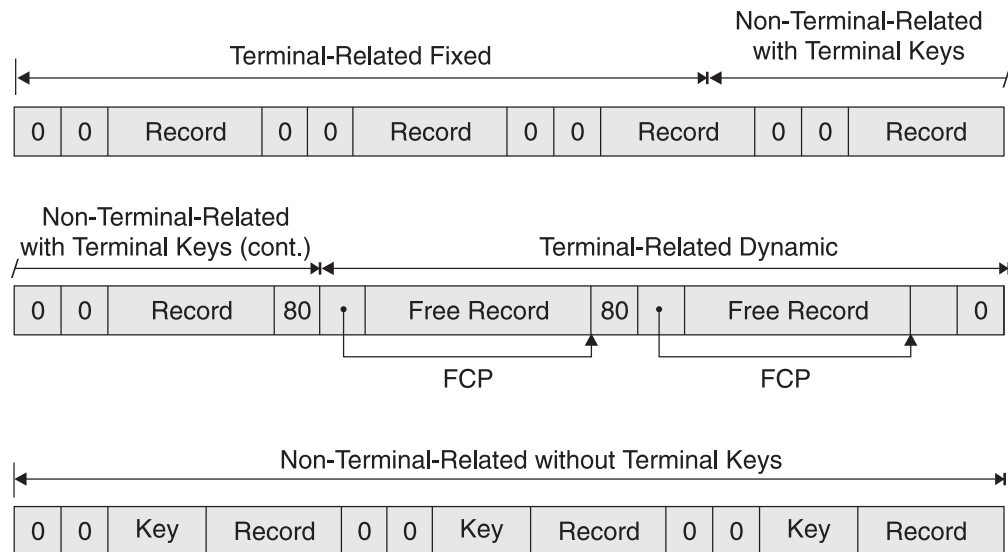


Figure 74. Sequence of the Four MSDB Organizations

Saving MSDBs for Restart

At system checkpoint, a copy of all MSDBs is written alternately to one of the MSDB checkpoint data sets—MSDBCP1 or MSDBCP2. During restart, the MSDBs are reloaded from the most recent copy on MSDBCP1 or MSDBCP2. During an emergency restart, the log is used to update the MSDB. During a normal restart, the operator can reload from MSDBINIT using the MSDBLOAD parameter on the restart command.

On a cold start (including /ERE CHKPT 0), MSDBs are loaded from the MSDBINIT data set.

DL/I Calls against an MSDB

All DL/I database calls, except those that specify “within parent”, are valid with MSDBs. Because an MSDB is a root-only database, a “within parent” call is meaningless. Additionally, the DL/I call, FLD, exists that is applicable to all MSDBs. The FLD call allows an application program to check and modify a single field in an MSDB segment.

Rules for Using an SSA

MSDB processing imposes the following restrictions on the use of an SSA (segment search argument):

- No boolean operator
- No command code

Even with the preceding restrictions, the result of a call to the database with no SSA, an unqualified SSA, or a qualified SSA remains the same as a call to the full-function database. For example, a retrieval call without an SSA returns the first record of the MSDB or the full-function database, depending on the environment in which you are working. The following list shows the type of compare or search technique used for a qualified SSA.

Type of Compare

Main Storage Databases (MSDBs)

- Sequence field: logical
- Non-sequence arithmetic field: arithmetic
- Non-sequence non-arithmetic: logical

Type of Search

- Sequence field: binary if operator is = or >=, otherwise sequential
- Non-sequence arithmetic field: sequential
- Non-sequence non-arithmetic: sequential

Insertion and Deletion of Segments

The terminal-related dynamic database accepts ISRT and DLET calls, and the other MSDB databases do not. Actual physical insertion and deletion of segments do not occur in the dynamic database. Rather, a segment is assigned to an LTERM from a pool of free segments by an ISRT call. The DLET call releases the segment back to the free segment pool.

Figure 75 on page 137 shows a layout of the four MSDBs and the control blocks and tables necessary to access them. The Extended Communications Node Table (ECNT) is located by a pointer from the Extended System Contents Directory (ESCD), which in turn is located by a pointer from the System Contents Directory (SCD). The ESCD contains first and last header pointers to the MSDB header queue. Each of the MSDB headers contains a pointer to the start of its respective database area.

Combination of Binary and Direct Access Methods

A combination access technique works against the MSDB on a DL/I call. The access technique combines a binary search and the direct access method. A binary search of the ECNT table attempts to match the table LTERM names to the LTERM name of the requesting terminal. When a match occurs, the application program accesses the segment of the desired database using a direct pointer in the ECNT table. Access to the non-terminal-related database segments without terminal keys is accomplished by a binary search technique only, without using the ECNT.

Figure 75 on page 137 shows the ECNT and MSDB storage layout.

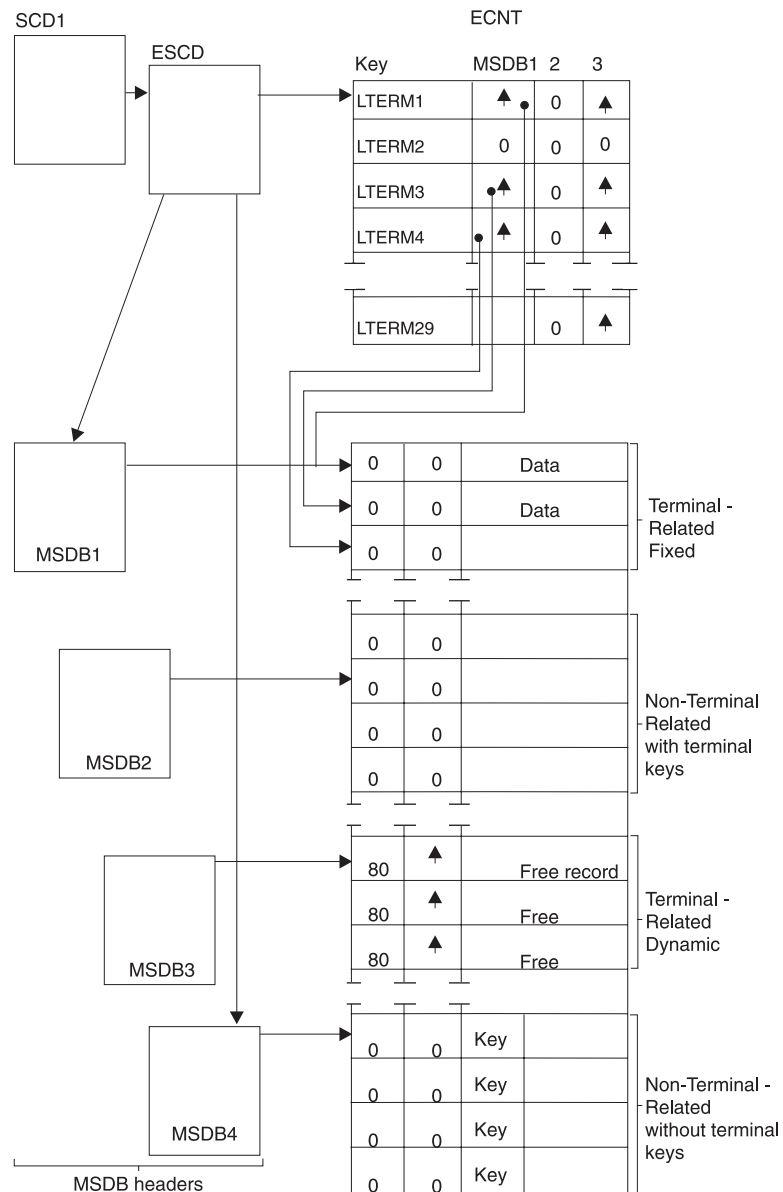


Figure 75. ECNT and MSDB Storage Layout

Position in an MSDB

Issuing a DL/I call causes a position pointer to fix on the current segment. The meaning of “next segment” depends on the key of the MSDB. The current segment in a non-terminal-related database without LTERM keys is the physical segment against which a call was issued. The next segment is the following physically adjacent segment after the current segment. The other three databases, using LTERM names as keys, have a current pointer fixed on a position in the ECNT table. Each entry in the table represents one LTERM name and segment pointers to every MSDB with which LTERM works. A zero entry indicates no association between an LTERM and an MSDB segment. If nonzero, the next segment is the next entry in the table. The zero entries are skipped until a nonzero entry is found.

The Field Call

The DL/I FLD call is available to MSDBs and DEDB. It allows for the operation on a field, rather than on an entire segment. Additionally, it allows conditional operation on a field.

Modification is done with the CHANGE form of the FLD call. The value of a field can be tested with the VERIFY form of the FLD call. These forms of the call allow an application program to test a field value before applying the change. If a VERIFY fails, all CHANGE requests in the same FLD call are denied. This call is described in *IMS Version 9: Application Programming: Database Manager*.

Call Sequence Results

The same call sequence against MSDBs and other IMS databases might bring different results. For parallel access to MSDB data, updates to MSDB records take place during sync point processing. Changes are not reflected in those records until the sync point is completed. For example, the sequence of calls GHU (Get-Hold-Unique), REPL (Replace), and GU (Get-Unique) for the same database record results in the same information in the I/O area for the GU call as that returned for the GHU.

The postponement of an updated database record to the point of commitment is also true of FLD/CHANGE calls, and affects FLD/VERIFY calls. You should watch for multiple FLD/VERIFY and FLD/CHANGE calls on the same field of the same segment. Such sequences can decrease performance because reprocessing results.

For terminal-related dynamic MSDBs, the following examples of call sequences do *not* have the same results as with other IMS databases or DEDBs:

- A GHU following an ISRT receives a 'segment not found' status code.
- An ISRT after a DLET receives a 'segment already exists' status code.
- No more than one ISRT or DLET is allowed for each MSDB in processing a transaction.

The preceding differences become more critical when transactions update or refer to both full function DL/I and MSDB data. Updates to full function DL/I databases and DEDBs are immediately available while MSDB changes are not. For example, if you issue a GHU and a REPL for a segment in an MSDB, then you issue another get call for the same segment in the same commit interval, the segment IMS returns to you is the “old” value, not the updated one.

If processing is not single mode, this difference can increase. In the case of multiple mode processing, the sync point processing is not invoked for every transaction. Your solution might be to ask for single mode processing when MSDB data is to be updated.

Another consideration for MSDB processing is that terminal-related MSDB segments can be updated only by transactions originating from the owners of the segment, the LTERMs. Programs that are non-transaction-driven BMPs can only update MSDBs that are declared as non-terminal-related.

Fast Path Virtual Storage Option

The Fast Path Virtual Storage Option (VSO) allows you to map data into virtual storage or a coupling facility structure. You can map one or more DEDB areas into virtual storage or a coupling facility structure by defining the DEDB areas as VSO areas.

For high-end performance applications with DEDBs, defining your DEDB areas as VSO allows you to realize the following performance improvements:

- Reduced read I/O

After an IMS and VSAM control interval (CI) has been brought into virtual storage, all subsequent I/O read requests read the data from virtual storage rather than from DASD.

- Decreased locking contention

For VSO DEDBs, locks are released after both of the following:

- Logging is complete for the second phase of an application synchronization (commit) point
- The data has been moved into virtual storage

For non-VSO DEDBs, locks are held at the VSAM CI-level and are released only after the updated data has been written to DASD.

- Fewer writes to the area data set

Updated data buffers are not immediately written to DASD; instead they are kept in the data space and written to DASD at system checkpoint or when a threshold is reached.

In all other respects, VSO DEDBs are the same as non-VSO DEDBs. Therefore, VSO DEDB areas are available for IMS DBCTL and LU 6.2 applications, as well as other IMS DB or IMS TM applications. Use the DBRC commands INIT.DBDS and CHANGE.DBDS to define VSO DEDB areas.

The virtual storage for VSO DEDB areas is housed differently depending on the share level assigned to the area. VSO DEDB areas with share levels of 0 and 1 are loaded into a z/OS data space. VSO DEDB areas with share levels of 2 and 3 are loaded into a coupling facility cache structure.

Coupling facility cache structures are defined by the system administrator and can accommodate either a single DEDB area or multiple DEDB areas. Cache structures that support multiple DEDB areas are called *multi-area structures*. For more information on multi-area structures, see *IMS Version 9: Administration Guide: System*.

Recommendation: Terminal-related MSDBs and non-terminal-related MSDBs with terminal-related keys are *not supported* in IMS Version 5 and later releases. Non-terminal-related MSDBs without terminal-related keys are still supported. Therefore, you should consider converting all your existing MSDBs to VSO DEDBs or non-VSO DEDBs.

Restrictions Using VSO DEDB Areas

VSO DEDB areas have the following restrictions in their use:

- VSO DEDB areas *must* be registered with DBRC.
- The maximum allowable size for either a z/OS data space or a coupling facility cache structure is 2 GB (2 147 483 648 bytes).

Fast Path Virtual Storage Option

The actual size available for a VSO area is the maximum size (2 GB) minus amounts used by z/OS (from 0 to 4 KB) and IMS Fast Path (approximately 100 KB). To see the size, usage, and other statistics for a VSO DEDB area, enter the /DISPLAY FPVIRTUAL command.

- The DEDB Area Data Set Compare utility (DBFUMMH0) does not support VSO DEDB areas.

Related Reading:

- See “Accessing a Data Space” on page 148 for more information on data spaces.
- See *IMS Version 9: Command Reference* for more information on the /DISPLAY commands.

Defining a VSO DEDB Area

All of the Virtual Storage Option (VSO) information for a DEDB is recorded in the RECON data set. Use the following parameters of the DBRC INIT.DBDS and CHANGE.DBDS commands to define your VSO DEDB Areas:

VSO Defines the area as a VSO area.

When a CI is read for the first time, it will be copied into a z/OS data space or a coupling facility structure. Data is read into a common buffer and is then copied into the data space or structure. Subsequent access to the data retrieves it from the data space or structure rather than from DASD.

CIs that are not read are not copied into the data space or structure.

All updates to the data are copied back to the data space or structure and any locks held are released. Updated CIs are periodically written back to DASD.

NOVSO

Defines the area as a non-VSO area. This is the default.

You can use NOVSO to define a DEDB as non-VSO or to turn off the VSO option for a given area. If the area is in virtual storage when it is redefined as NOVSO, the area must be stopped (/STOP AREA or /DBR AREA) or removed from virtual storage (/VUNLOAD) for the change to take effect.

PRELOAD

For VSO areas, this preloads the area into the data space or coupling facility structure when the VSO area is opened. This keyword implies the PREOPEN keyword, thus if PRELOAD is specified, then PREOPEN does not have to be specified.

The root addressable portion and the independent overflow portion of an area are loaded into the data space or coupling facility structure at control region initialization or during /START AREA processing. Data is then read from the data space or coupling facility structure to a common buffer. Updates are copied back to the data space or coupling facility structure and any locks are released. Updated CIs are periodically written back to DASD.

NOPREL

Defines the area as load-on-demand. For VSO DEDBs areas, as CIs are read from the data set, they are copied to the data space or coupling facility structure. This is the default.

To define an area with NOPREL gives you the ability to deactivate the preload processing. The area is not preloaded into the data space or coupling facility structure the next time that it is opened.

If you specify NOPREL, and you want the area to be preopened, you must separately specify PREOPEN for the area.

CFSTR1

Defines the name of the cache structure in the primary coupling facility. Cache structure names must follow z/OS coupling facility naming conventions. CFSTR1 uses the name of the DEDB area as its default. This parameter is valid only for VSO DEDB areas that are defined with SHARELVL(2I3).

Related Reading: For detailed information on coupling facility naming, see “Coupling Facility Structure Naming Convention” on page 144.

CFSTR2

Defines the secondary coupling facility cache structure name when you use IMS-managed duplexing of structures. The cache structure name must follow z/OS coupling facility naming conventions. CFSTR2 does not provide a default name. This parameter is valid only for VSO areas of DEDBs that are defined with SHARELVL(2I3) and that are single-area structures. This parameter cannot be used with multi-area structures, which use system-managed duplexing.

Related Reading:

- For detailed information on coupling facility naming, see “Coupling Facility Structure Naming Convention” on page 144.
- For more information on multi-area structures, see *IMS Version 9: Administration Guide: System*.

MAS

Defines a VSO DEDB area as using a multi-area structure as opposed to a single-area structure.

Related Reading: For more information on multi-area structures, see *IMS Version 9: Administration Guide: System*.

NOMAS

Defines a VSO DEDB area as using a single-area cache structure as opposed to a multi-area structure. NOMAS is the default.

LKASID

Indicates that buffer lookaside is to be performed on read requests for this area. For VSO DEDB areas that use a multi-area structure, lookaside can also be specified using the DFSVSMxx PROCLIB member. If there is a discrepancy between the specifications in DBRC and those in DFSVSMxx, the specifications in DFSVSMxx are used.

Related Reading: For additional information on defining private buffer pools, see “Defining a Private Buffer Pool Using the DFSVSMxx IMS.PROCLIB Member” on page 146.

NOLKASID

Indicates that buffer lookaside is not to be performed on read requests for this area.

Related Reading: For additional information on defining private buffer pools, see “Defining a Private Buffer Pool Using the DFSVSMxx IMS.PROCLIB Member” on page 146.

VSO DEDB Areas and the PREOPEN and NOPREO Keywords

The PREOPEN and NOPREO keywords of DBRC’s INIT.DBDS and CHANGE.DBDS commands apply to both VSO DEDB areas and non-VSO DEDB areas.

Fast Path Virtual Storage Option

When a NOPREO area is also defined as shared VSO with a share level of 2 or 3, you can open the area with the /START AREA command. This connects the area to the VSO structures.

You can use the DBRC commands to define your VSO DEDB areas at any time; it is not necessary that IMS be active. The keywords specified on these DBRC commands go into effect at two different points in Fast Path processing:

- Control region startup

After the initial checkpoint following control region initialization, DBRC provides a list of areas with any of the VSO options (VSO, NOVSO, PRELOAD, and NOPREL) or either of the PREOPEN or NOPREO options. The options are then maintained by IMS Fast Path.

- Command processing

When you use a /START AREA command, DBRC provides the VSO options or PREOPEN/NOPREO options for the area. If the area needs to be preopened or preloaded, it is done at this time.

When you use a /STOP AREA command, any necessary VSO processing is performed.

Related Reading: See *IMS Version 9: Command Reference* for details on start and stop processing.

Sharing of VSO DEDB Areas

Sharing of VSO DEDB areas allows multiple IMS systems to concurrently read and update the same VSO DEDB area. The three main participants are the coupling facility hardware, the coupling facility policy software, and the XES and z/OS services.

The coupling facility hardware provides high-performance, random-access shared storage in which IMS systems can share data in a sysplex environment. The shared storage area in the coupling facility is divided into sections, called *structures*. For VSO DEDB data, the structure type used is called a *cache structure*, as opposed to a list structure or a lock structure. The cache structure is designed for high-performance read reference reuse and deferred write of modified data. The coupling facility and structures are defined in a common z/OS data set, the *couple data set* (COUPLExx).

The coupling facility policy software and its cache structure services provide interfaces and services to z/OS that allow sharing of VSO DEDB data in shared storage. Shared storage controls VSO DEDB reads and writes:

- A read of a VSO CI brings the CI into the coupling facility from DASD.
- A write of an updated VSO CI copies the CI to the coupling facility from main storage, and marks it as changed.
- Changed CI data is periodically written back to DASD.

The XES and z/OS services provide a way of manipulating the data within the cache structures. They provide high performance, data integrity, and data consistency for multiple IMS systems sharing data.

The Coupling Facility and Shared Storage

In the coupling facility shared storage, a cache structure can represent one or multiple VSO DEDB areas; however, any given VSO DEDB area can be represented by only one cache structure. Cache structures are not persistent. That is, they are deleted after the last IMS system disconnects from the coupling facility.

Duplexing Structures

Duplexed structures are duplicate structures for the same area. Duplexing allows you to have dual structure support for your VSO DEDB areas, which helps to ensure the availability and recoverability of your data.

Structure duplexing can be either IMS-managed or system-managed. With IMS-managed duplexing, you must define both the primary and the secondary structures in DBRC and in the z/OS coupling facility resource management (CFRM) policy. When you use system-managed duplexing, you have to define only the primary structure. The duplexing operation is transparent to you, except that you need to request duplex mode in your CFRM policy and allocate additional resources for a secondary structure instance.

VSO multi-area structures require the use of system-managed duplexing.

Related Reading: For information about enabling and initiating system-managed duplexing, see the chapter on data sharing in *IMS Version 9: Administration Guide: System*.

Automatic Altering of Structure Size

z/OS can automatically expand or contract the size of a VSO structure in the coupling facility if it needs storage space. Enabling this function for preloaded VSO DEDBs can prevent wasted space; however, you must be careful with this function when VSO DEDBs are loaded on demand.

If you have dual structures, IMS systems below Version 8 cannot connect to structures with different sizes.

Recommendation: If you preload your shared VSO DEDB areas, do not use the automatic alter function. The automatic alter function might reclaim any unchanged data in the cache structure.

To ensure correct sizing and that the automatic alter function is disabled for your preloaded shared VSO DEDB areas, specify the following parameters in the area's CFRM policy:

INITSIZE(x)

Specifies the initial amount of space to be allocated for the structure in the coupling facility. You can estimate this size using the CFSIZER tool available online.

SIZE(y)

Specifies the maximum amount of space to be allocated for the structure in the coupling facility.

ALLOWAUTOALT(NO)

Specifies whether to allow system-initiated alters (automatic alter) for this structure. For preloaded shared VSO DEDB areas, you should specify ALLOWAUTOALT(NO).

If the size specified by INITSIZE is too small, IMS will alter the size of the structure to the needed value as calculated by IXLCSF, up to the size specified in the SIZE parameter.

Related Reading: For information on the CFRM parameters that enable automatic altering of structures, see:

- The chapter on data sharing in *IMS Version 9: Administration Guide: System*

- *z/OS V1R4: MVS Setting Up a Sysplex*

System-Managed Rebuild

You can reconfigure a coupling facility while keeping all VSO structures online by copying the structures to another coupling facility. There is no change to the VSO definition.

Related Reading: For information on enabling and allocating a system-managed rebuild, allocating and populating a new structure, and managing the coupling facility, see the chapter on data sharing in *IMS Version 9: Administration Guide: System*.

Private Buffer Pools

IMS provides special private buffer pools for Shared VSO areas. Each pool can be associated with an area, a DBD, or a specific group of areas. These private buffer pools are only used for Shared VSO data. Using these private buffer pools, the customer can request buffer lookaside for the data. The keywords LKASID or NOLKASID, when specified on the DBRC commands INIT.DBDS or CHANGE.DBDS, indicate whether to use this lookaside capability or not.

Defining a VSO Cache Structure Name

The system programmer defines all coupling facility structures, including VSO cache structures, in the CFRM policy definition. In this policy definition, VSO structures are defined as cache structures, as opposed to list structures (used by shared queues) or lock structures (used by IRLM).

Coupling Facility Structure Naming Convention

The structure name is 16 characters long, padded on the right with blanks if necessary. It can contain any of the following, but must begin with an uppercase, alphabetic character:

- Uppercase alphabetic characters
- Numeric characters
- Special characters (\$, @, and #)
- Underscore (_)

IBM names begin with:

- 'SYS'
- Letters 'A' through 'I' (uppercase)
- An IBM component prefix

Examples of Defining Coupling Facility Structures

Figure 76 on page 145 shows how to define two structures in separate coupling facilities.

```

//UPDATE EXEC PGM=IXCL2FDA
//SYSPRINT DD SYSOUT=A
//*
/* THE FOLLOWING SYSIN WILL UPDATE THE POLICY DATA IN THE COUPLE
/* DATASET FOR CFRM (COUPLING FACILITY RESOURCE MANAGEMENT)
/*
//SYSIN DD *
UPDATE DSN(IMS.DSHR.PRIME.FUNC) VOLSER(DSHR03)

DEFINE POLICY(POLICY1)

    DEFINE CF(FACIL01)
        ND(123456)
        SIDE(0)
        ID(01)
        DUMPSPACE(2000)

    DEFINE CF(FACIL02)
        ND(123456)
        SIDE(1)
        ID(02)
        DUMPSPACE(2000)

    DEFINE STR(LIST01)
        SIZE(1000)
        PREFLIST(FACIL01,FACIL02)
        EXCLLIST(CACHE01)

    DEFINE STR(CACHE01)
        SIZE(1000)
        PREFLIST(FACIL02,FACIL01)
        EXCLLIST(LIST01)
/*

```

Figure 76. Example of Updating a Policy with New Structures

In the example, the programmer defined one list structure (LIST01) and one cache structure (CACHE01).

Attention: When defining a cache structure to DBRC, ensure that the name is identical to the name used in the CFRM policy (see “Registering a Cache Structure Name with DBRC”).

Registering a Cache Structure Name with DBRC

When you define DEDB areas to DBRC, use the same structure names defined in the CFRM policy to specify the structures that each DEDB area will use. The DEDB area definitions and the corresponding structure names are then stored in the RECON data set. The structure names are entered in either the CFSTR1 or CFSTR2 parameter of the INIT.DBDS or CHANGE.DBDS command. For more information on defining DEDB areas, see “Defining a VSO DEDB Area” on page 140.

Restriction: The CFSTR2 parameter is not supported by multi-area structures. If you specify both CFSTR2 and MAS in INIT.DBDS, or use CHANGE.DBDS to apply CFSTR2 to DEDB area already defined by MAS, IMS will reject the DBRC command with either a DSP0141I or DSP0144I error message.

Figure 77 on page 146 registers structure name TSTDDBAR1.


```
INIT,DBDS DBD(DEDBFR01) AREA(DEDBAR1) VSO PRELOAD CFSTR1(TSTDDBAR1)
```

Figure 77. Defining a VSO Area Coupling Facility Structure Name in DBRC

Defining a Private Buffer Pool Using the DFSVSMxx IMS.PROCLIB Member

Define a private buffer pool using the following format:

```
DEDB=(poolname,size,pbuf,sbuf,maxbuf,LKASID|NOLKASID,dbname)
```

where:

| | |
|-----------------|---|
| poolname | 8 character name of the pool. Used in displays and reports. |
| size | The buffer size of the pool. All the standard DEDB-supported buffer sizes are supported. |
| pbuf | The primary buffer allocation. The first allocation receives this number of buffers. Maximum value is 99999. |
| sbuf | The secondary buffer allocation. If the primary allocation starts to run low, another allocation of buffers is made. This amount indicates the secondary allocation amount. Maximum value is 99999. |
| maxbuf | The maximum number of buffers allowed for this pool. It is a combination of PBUF plus some iteration of SBUF. Maximum value is 99999. |

LKASIDINOLKASID

Indicates whether this pool is to be used as a local cache with buffer lookaside capability. This value is cross-checked with the DBRC specification of LKASID to determine which pool the area will use. If there is an inconsistency between the DEDB statement and DBRC, the DBRC value takes precedence.

| | |
|---------------|---|
| dbname | Association of the pool to a specific area or DBD. If the dbname is an area name, then the pool is used only by that area. If the dbname specifies a DBD name, the pool is used by all areas in that DBD. The dbname is first checked for an area name then for a DBD name. |
|---------------|---|

Figure 78 shows examples of how to define a private buffer pool.

```
DEDB=(POOL1,512,400,50,800,LKASID)
DEDB=(POOL2,8196,100,20,400,NOLKASID)
```

Figure 78. Examples of Defining Private Buffer Pools

In this example, 2 private buffer pools are defined:

1. The first pool has a buffer size of 512, with an initial allocation of 400 buffers, increasing by 50, as needed, to a maximum of 800. This pool will be used as a local cache, and buffer lookaside will be performed for areas that share this pool.
2. The second pool has a buffer size of 8K, with an initial allocation of 100 buffers, increasing by 20, as needed, to a maximum of 400. This pool will be used in the same fashion as the common buffer pool. There will be no lookaside performed.

If the customer does not define a private buffer pool, the default parameter values are calculated as follows:

DEDB=(poolname,XXX,64,16,512)

where:

- XXX is the CI size of the area to be opened.
- The initial buffer allocation is 64.
- The secondary allocation is 16.
- The maximum number of buffers for the pool is 512.
- The LKASID option is specified if it is specified in DBRC for the area.

Defining a Private Buffer Pool for a Multi-Area Structure

You can define private buffer pools for multi-area structure using the DEDBMAS= keyword in the DFSVSMxx PROCLIB member. The format is as follows:

DEDBMAS=(poolname,cisize,pbuf,sbuf,maxbuf,LKASID|NOLKASID,strname)

Except for the following parameters, the parameters for DEDBMAS are the same as those in the DFSVSMxx DEDB= keyword:

| | |
|----------------|---|
| cisize | The control interval size of the area. All areas that share a multi-area structure must have the same control interval size. If there is a discrepancy between the control interval size of the area used in creating the structure and the control interval size of the area attempting to share the structure, the open process for the area attempting to share the structure fails. |
| strname | The required 1- to 16-character name of the primary coupling facility structure. The installation must have defined the structure in the CFRM administrative policy. The structure name must follow the naming conventions of the CFRM. If the name has fewer than 16 characters, the system pads the name with blanks. The valid characters are A–Z, 0–9, and the characters \$, &, #, and _. Names must be uppercase and start with alphabetic character. |
| | Restriction: Do not begin structure names with the letters A-I, or the character string SYS. IBM reserves these characters for its structures. |

Acquiring and Accessing Data Spaces for VSO DEDB Areas

IMS allocates data spaces to accommodate VSO DEDB areas. When a VSO DEDB area CI is preloaded or read for the first time, it is copied into a data space (or a coupling facility structure). Subsequent access to the data retrieves it from the data space rather than from DASD.

Acquiring a Data Space

IMS acquires data spaces for VSO areas when the VSO areas first open, but not before. The maximum size of any VSO area data space is two gigabytes. Data spaces for preloaded VSO areas use the z/OS DREF (disabled reference) option. Data spaces for non-preloaded VSO areas do not use the DREF option.

DREF data spaces use a combination of central storage and expanded storage, but no auxiliary storage. Data spaces without the DREF option use central storage, expanded storage, and auxiliary storage, if auxiliary storage is available.

IMS acquires additional data spaces for VSO areas, both with DREF and without, as needed.

Accessing a Data Space

During IMS control region initialization, IMS calls DBRC to request a list of all the areas that are defined as VSO. This list includes the PREOPEN or PRELOAD status of each VSO area. If VSO areas exist, IMS acquires the appropriate data spaces. Then IMS opens all areas defined with PREOPEN and opens and loads areas defined with PRELOAD. During a normal or emergency restart, the opening and loading of areas might occur after control region initialization, if you have changed the specifications of the FPOPN parameter in the IMS procedure.

IMS assigns areas to data spaces using a “first fit” algorithm. The entire root addressable portion of an area (including independent overflow) resides in the data space. The sequential dependent portion does not reside in the data space.

The amount of space needed for an area in a data space is $(CI\ size) \times (number\ of\ CIs\ per\ UOW) \times ((number\ of\ UOWs\ in\ root\ addressable\ portion) + (number\ of\ UOWs\ in\ independent\ overflow\ portion))$ rounded to the next 4 KB.

Expressed in terms of the parameters of the DBDGEN AREA statement, this formula is $(SIZE\ parameter\ value) \times (UOW\ parameter\ value) \times (ROOT\ parameter\ value)$ rounded to the next 4 KB.

The actual amount of space in a data space available for an area (or areas) is two gigabytes (524,288 blocks, 4 KB each) minus an amount reserved by z/OS (from 0 to 4 KB) minus an amount used by IMS Fast Path (approximately 100 KB). You can use the /DISPLAY FPVIRTUAL command to determine the actual storage usage of a particular area.

Related Reading: For sample output from this command, see *IMS Version 9: Command Reference*.

Resource Control and Locking

Using VSO can reduce the number and duration of DEDB resource locking contentions by managing DEDB resource requests on a segment level and holding locks only until updated segments are returned to the data space. Segment-level resource control and locking applies only to Get and Replace calls.

Without VSO, the VSAM CI (physical block) is the smallest available resource for DEDB resource request management and locking. If there is an update to any part of the CI, the lock is held until the whole CI is rewritten to DASD. No other requester is allowed access to any part of the CI until the first requester's lock is released.

With VSO, the database segment is the smallest available resource for DEDB resource request management and locking. Segment-level locking is available only for the root segment of a DEDB with a root-only structure, and when that root segment is a fixed-length segment. If processing options R or G are specified in the calling PCB, IMS can manage and control DEDB resource requests and serialize change at the segment level; for other processing options, IMS maintains VSAM CI locks. Segment locks are held only until the segment updates are applied to the CI in the data space. Other requesters for different segments in the same CI are allowed concurrent access.

A VSO DEDB resource request for a segment causes the entire CI to be copied into a common buffer. VSO manages the segment request at a level of control consistent with the request and its access intent. VSO also manages access to the

CI that contains the segment but at the share level in all cases. A different user's subsequent request for a segment in the same CI accesses the image of the CI already in the buffer.

Updates to the data are applied directly to the CI in the buffer at the time of the update. Segment-level resource control and serialization provide integrity among multiple requesters. After an updated segment is committed and applied to the copy of the CI in the data space, other requesters are allowed access to the updated segment from the copy of the CI in the buffer.

If after a segment change the requester's updates are not committed for any reason, VSO copies the unchanged image of the segment from the data space to the CI in the buffer. VSO does not allow other requesters to access the segment until VSO completes the process of removing the uncommitted and cancelled updates. Locking at the segment level is not supported for shared VSO areas. Only CI locking is supported.

When a compression routine is defined on the root segment of a DEDB with a root-only structure, and when that root segment is a fixed-length segment, its length becomes variable after being compressed. Replacing a compressed segment then requires a delete and an insert. In this case, segment level control and locking is not available.

Preopen Areas and VSO Areas in a Data Sharing Environment

A VSO can be registered with any of the following share levels:

SHARELVL(0)

Exclusive access: in a data sharing environment, any SHARELVL(0) area with the PREOPEN option (including VSO PREOPEN and VSO PRELOAD) is opened by the first IMS system to complete its control region initialization. IMS will not attempt to preopen the area for any other IMS.

SHARELVL(1)

One updater, many readers: in a data sharing environment, a SHARELVL(1) area with the PREOPEN option is preopened by all sharing IMS systems. The first IMS system to complete its control region initialization has update authorization; all others have read authorization.

If the SHARELVL(1) area is a VSO area, it is allocated to a data space by any IMS that opens the area. If the area is defined as VSO PREOPEN or VSO PRELOAD, it is allocated to a data space by all sharing IMS systems.

If the area is defined as VSO NOPREO NOPREL, it is allocated to a data space by all IMS systems, as each opens the area. The first IMS to access the area has update authorization; all others have read authorization.

SHARELVL(2)

Block-level sharing: a SHARELVL(2) area with at least one coupling facility structure name (CFSTR1) defined is shared at the block or control interval (CI) level within the scope of a single IRLM. Multiple IMS systems can be authorized for update or read processing if they are using the same IRLM.

SHARELVL(3)

Block-level sharing: a SHARELVL(3) area with at least one coupling facility structure name (CFSTR1) defined is shared at the block or control interval (CI) level within the scope of multiple IRLMs. Multiple IMS systems can be authorized for nonexclusive access.

Fast Path Virtual Storage Option

Attention: Be careful when registering a VSO area as SHARELVL(1). Those systems that receive read-only authorization never see the updates made by the read/write system because all reads come from the data space (not from DASD, where updates are eventually written).

Input/Output Processing With VSO

This topic describes how IMS uses buffers, data spaces, and DASD in response to read and update requests.

Input Processing

When an application program issues a read request to a VSO area, IMS checks to see if the data is in the data space. If the data is in the data space, it is copied from the data space into a common buffer and passed back to the application. If the data is not in the data space, IMS reads the CI from the area data set on DASD into a common buffer, copies the data into the data space, and passes the data back to the application.

For SHARELVL(2|3) VSO areas, Fast Path uses private buffer pools. Buffer lookaside is an option for these buffer pools. When a read request is issued against a SHARELVL(2|3) VSO area using a lookaside pool, a check is made to see if the requested data is in the pool. If the data is in the pool, a validity check to XES is made. If the data is valid, it is passed back to the application from the local buffer. If the data is not found in the local buffer pool or XES indicates that the data in the pool is not valid, the data is read from the coupling facility structure and passed to the application. When the buffer pool specifies the no-lookaside option, every request for data goes to the coupling facility.

For those areas that are defined as load-on-demand (using the VSO and NOPREL options), the first access to the CI is from DASD. The data is copied to the data space and then subsequent reads for this CI retrieve the data from the data space rather than from DASD. For those areas that are defined using the VSO and PRELOAD options, all access to CIs comes from the data space.

Whether the data comes from DASD or from the data space is transparent to the processing done by application programs.

Output Processing

During phase 1 of synchronization point processing VSO data is treated the same as non-VSO data. The use of VSO is transparent to logging.

During phase 2 of the synchronization point processing VSO and non-VSO data are treated differently. For VSO data, the updated data is copied to the data space, the lock is released and the buffer is returned to the available queue. The relative byte address (RBA) of the updated CI is maintained in a bitmap. If the RBA is already in the bitmap from a previous update, only one copy of the RBA is kept. At interval timer, the updated CIs are written to DASD. This batching of updates reduces the amount of output processing for CIs that are frequently updated. While the updates are being written to DASD, they are still available for application programs to read or update because copies of the data are made within the data space just before it is written.

For SHARELVL(2|3) VSO areas, the output thread process is used to write updated CIs to the coupling facility structures. When the write is complete, the lock is released. XES maintains the updated status of the data in the directory entry for the CI.

The PRELOAD Option

The loading of one area takes place asynchronously with the loading of any others. The loading of an area is (or can be) concurrent with an application program's accesses to that area. If the CI requested by the application program has been loaded into the data space, it is retrieved from the data space. If the requested CI has not yet been loaded into the data space, it is obtained from DASD and UOW locking is used to maintain data integrity.

The preload process for SHARELVL(2|3) VSO areas is similar to that of SHARELVL(0|1). Multiple preloads can be run concurrently, and also concurrent with application processing. The locking, however, is different. SHARELVL(2|3) Areas that are loaded into coupling facility structures use CI locking instead of UOW locking. The load process into the coupling facility is done one CI at a time.

If a read error occurs during preloading, an error message flags the error, but the preload process continues. If a subsequent application program call accesses a CI that was not loaded into the data space due to a read error, the CI request goes out to DASD. If the read error occurs again, the application program receives an "AO" status code, just as with non-VSO applications. If instead the access to DASD is successful this time, the CI is loaded into the data space.

I/O Error Processing

Using VSO increases the availability of data when write errors occur. When a CI for a VSO area has been put into a data space, the CI is available from that data space as long as IMS is active, even if a write error occurs when an update to the CI is being written to DASD.

Write Errors: When a write error occurs, IMS create an error queue element (EQE) for the CI in error. For VSO areas, all read requests are satisfied by reading the data from the data space. Therefore, as long as the area continues to reside in the data space, the CI that had the write error continues to be available. When the area is removed from the data space, the CI is no longer available and any request for the CI receives an "AO" status code.

Read Errors: For VSO areas, the first access to a CI causes it to be read from DASD and copied into the data space. From then on, all read requests are satisfied from the data space. If there is a read error from the data space, z/OS abends.

For VSO areas that have been defined with the PRELOAD option, the data is preloaded into the data space; therefore, all read requests are satisfied from the data space.

Related Reading: See "The PRELOAD Option" for a discussion of read error handling during the preload process.

To provide for additional availability, SHARELVL(2|3) VSO areas support multiple structures per area. If a read error occurs from one of the structures, the read is attempted from the second structure. If there is only one structure defined and a read error occurs, an AO status code is returned to the application.

There is a maximum of three read errors allowed from a structure. When the maximum is reached and there is only one structure defined, the area is stopped and the structure is disconnected.

When the maximum is reached and there are two structures defined, the structure in error is disconnected. The one remaining structure is used. If a write error to a

Fast Path Virtual Storage Option

structure occurs, the CI in error is deleted from the structure and written to DASD. The delete of the CI is done from the sharing partners. If none of the sharers can delete the CI from the structure, an EQE is generated and the CI is deactivated. A maximum of three write errors are allowed to a structure. If there are two structures defined and one of them reaches the maximum allowed, it is disconnected.

Checkpoint Processing

During a system checkpoint, all of the VSO area updates that are in the data space are written to DASD. All of the updated CIs in the CF structures are also written to DASD. Only CIs that have been updated are written. Also, all updates that are in progress are allowed to complete before checkpoint processing continues.

VSO Options Across IMS Restart

For all types of IMS restart except XRF takeover (cold start, warm start, emergency restart, COLDBASE, COLDCOMM and COLDSYS emergency restart), the VSO options in effect after restart are those defined to DBRC. In the case of the XRF takeover, the VSO options in effect after the takeover are the same as those in effect for the active IMS prior to the failure that caused the XRF takeover.

Emergency Restart Processing

Recovery of VSO areas across IMS or z/OS failures is similar to recovery of existing non-VSO areas. IMS examines the log records, from a previous system checkpoint to the end of the log, to determine if there are any committed updates that were not written to DASD before the failure. If any such committed updates are found, IMS will REDO them (apply the update to the CI and write the updated CI to DASD). Because VSO updates are batched together during normal processing, VSO areas are likely to require more REDO processing than non-VSO areas.

During emergency restart log processing, IMS tracks VSO area updates differently depending on the share level of the VSO area. For share levels 0 and 1, IMS uses data spaces to track VSO area updates. For share levels 2 and 3, IMS uses a buffer in memory to track VSO area updates.

IMS also obtains a single non-DREF data space which it releases at the end of restart. If restart log processing is unable to get the data space or main storage resources it needs to perform VSO REDO processing, the area is stopped and marked as “recovery needed”.

By default, at the end of emergency restart, IMS opens areas defined with the PREOPEN or PRELOAD options. IMS then loads areas with the PRELOAD option into a data space or coupling facility structure. You can alter this behavior by using the FPOPN keyword of the IMS procedure to have IMS restore all VSO DEDB areas to their open or closed state at the time of the failure.

Related Reading: For more information on specifying how IMS reopens DEDB areas during an emergency restart, see “Reopening DEDB Areas During an Emergency Restart” on page 115.

VSO areas without the PREOPEN or PRELOAD options are assigned to a data space during the first access following emergency restart.

After an emergency restart, the VSO options and PREOPENINOPREO options in effect for an area are those that are defined to DBRC, which may not match those in effect at the time of the failure. For example, a non-shared VSO area removed

from virtual storage by the /VUNLOAD command before the failure, is restored to the data space after the emergency restart. For shared VSO areas, the area remains unloaded until the next /STA AREA command is issued for it.

VSO Options with XRF

During the tracking and takeover phases on the alternate IMS, log records are processed in the same manner as during active IMS emergency restart (from a previous active system checkpoint to the end of the log). The alternate IMS uses the log records to determine which areas have committed updates that were not written to DASD before the failure of the active IMS. If any such committed updates are found, the alternate will REDO them, following the same process as for active IMS emergency restart.

Related Reading: See “Emergency Restart Processing” on page 152 for information on restart and REDO.

During tracking, the alternate uses data spaces to track VSO area updates: in addition to the data space resources used for VSO areas, the alternate obtains a single non-DREF data space which it releases at the end of takeover. If XRF tracking or takeover is unable to get the data space or main storage resources it needs to perform VSO REDO processing, the area is stopped and marked “recovery needed”.

Following an XRF takeover, areas that were open or in the data space remain open or in the data space. The VSO options and PREOPEN|NOPRE0 options that were in effect for the active IMS before the takeover remain in effect on the alternate (the new active) after the takeover. Note that these options may not match those defined to DBRC. For example, a VSO area removed from virtual storage by the /VUNLOAD command before the takeover is *not* restored to the data space after the takeover.

VSO areas defined with the preload option are preloaded at the end of the XRF takeover. In most cases, dependent regions can access the area before preloading begins, but until preloading completes, some area read requests may have to be retrieved from DASD.

Fast Path Synchronization Points

MSDBs and DEDBs are not updated during application program processing, but the updates are kept in buffers until a sync point. Output messages are not sent until the message response is logged. The Fast Path sync point is defined as the next GU call for a message-driven program, or a SYNC or CHKP call for a BMP using Fast Path facilities. Sync point processing occurs in two phases.

Phase 1 - Build Log Record

DEDB updates and verified MSDB records are written in system log records. All DEDB updates for the current sync point are chained together as a series of log records. Resource contentions, deadlocks, out-of-space conditions, and MSDB verify failures are discovered here.

Phase 2 - Write Record to System Log

Database and message records are written to the IMS system log. After logging, MSDB records are updated, the DEDB updates begin, and messages are sent to the terminals. DEDB updates are applied with a type of asynchronous processing

Fast Path Synchronization Points

called an output thread. Until the DEDB changes are made, any program that tries to access unwritten segments is put in a wait state.

If, during application processing, a Fast Path program issues a call to a database other than MSDB or DEDB, or to an alternate PCB, the processing is serialized with full function events. This can affect the performance of the Fast Path program. In the case of a BMP or MPP making a call to a Fast Path database, the Fast Path resources are held, and the throughput for Fast Path programs needing these resources can be affected.

Managing I/O Errors and Long Wait Times

When a database I/O error occurs in single area data sets (ADS), IMS copies the buffer contents of the error control interval (CI) to a virtual buffer. A subsequent DL/I request causes the error CI to be read back into the buffer pool. The write error information and buffers are maintained across restarts, allowing recovery to be deferred to a convenient time. I/O error retry is automatically performed at database close time and at system checkpoint. If the retry is successful, the error condition no longer exists and recovery is not needed.

Multiple Area Data Sets I/O Timing (MADSIOT) helps you avoid the excessively long wait times (also known as a *long busy*) that can occur while a RAMAC® disk array performs internal recovery processing.

Restriction: MADSIOT applies only to multiple area data sets (MADS). For single area data sets (ADS), IMS treats the long busy condition as a permanent I/O error handled by the Fast Path I/O toleration function. The MADSIOT function works only on a system that supports the long busy state.

To invoke MADSIOT, you must define the MADSIOT keyword on the DFSVSMxx PROCLIB member. The /STA MADSIOT and /DIS AREA MADSIOT commands serve to start and monitor the MADSIOT function.

Additionally, MADSIOT requires the use of a Coupling Facility (CFLEVEL=1 or later) list structure in a sysplex environment. MADSIOT uses this Coupling Facility to store information required for DB recovery. You must use the CFRM policy to define the list structure name, size, attributes, and location.

Table 15 shows the required CFRM list structure storage sizes when CFLEVEL=12 and the number of changed CIs is 1 000, 5 000, 20 000, and 30 000. The storage sizes differ at different CFLEVELs.

Table 15. Required CFRM List Structure Storage Sizes for CFLEVEL=12

| Altered number of CIs (entrynum) | Required Storage Size (listheadernum=50) |
|-------------------------------------|--|
| 1 000 | 1 792 KB |
| 5 000 | 3 584 KB |
| 20 000 | 11 008 KB |
| 30 000 | 15 616 KB |

You can estimate CFRM list structure storage sizes tailored to your installation using an online tool: the IBM S/390 Coupling Facility Structure Sizer Tool (CFSizer). CFSizer is available for you to use at the following Web site:

www-1.ibm.com/servers/eserver/zseries/cfsizer/, or search for “CFSizer” at the IBM Web site, www.ibm.com.

Related Reading:

- For additional information on the MADSIOT keyword, see the topic on the DFSVSMxx PROCLIB member in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.
- For an example of defining CFRM policies, see the *IMS Version 9: Common Queue Server Guide and Reference*.
- For information on the /STA MADSIOT and /DIS AREA MADSIOT commands, see the *IMS Version 9: Command Reference*.

Registering Fast Path Databases in DBRC

Although databases need not be registered in DBRC in order for the error handling to work, registration is highly recommended. If an error occurs on a database not registered and the system stops, the database could be damaged if the system is restarted and a /DBR command is not issued prior to accessing the database. The restart causes the error buffers to be restored as they were when the system stopped. If the same block had been updated during the batch run, the batch update would be overlaid.

Chapter 8. Choosing Optional Database Functions

After you have determined the type of database that best suits your application's processing requirements, you are ready to determine which additional IMS functions you need to use.

This chapter explains the following functions and describes when and how to use them:

- “Logical Relationships”
- “Secondary Indexes” on page 192
- “Variable-Length Segments” on page 215
- “Segment Edit/Compression Exit Routine” on page 218
- “Data Capture Exit Routines” on page 221
- “Field-Level Sensitivity” on page 226
- “Multiple Data Set Groups” on page 235
- “VSAM KSDS CI Reclaim For Full-Function Databases” on page 241
- “HALDB Single Partition Processing” on page 242
- “Integrated HALDB Online Reorganization Function” on page 244
- “Storing XML Data in IMS Databases” on page 244

Notes:

1. These functions do not apply to GSAM, MSDB, HSAM, and SHSAM databases.
2. Only the variable-length segment function, the Segment Edit/Compression exit routine, and the Data Capture exit routine apply to DEDBs.

Logical Relationships

The following database types support logical relationships:

- HISAM
- SHISAM
- HDAM
- PHDAM
- HIDAM
- PHIDAM

Logical relationships resolve conflicts in the way application programs need to view segments in the database. With logical relationships, application programs can access:

- Segment types in an order other than the one defined by the hierarchy
- A data structure that contains segments from more than one physical database.

An alternative to using logical relationships to resolve the different needs of applications is to create separate databases or carry duplicate data in a single database. However, in both cases this creates duplicate data. Avoid duplicate data because:

- Extra maintenance is required when duplicate data exists because both sets of data must be kept up to date. In addition, updates must be done simultaneously to maintain data consistency.
- Extra space is required on DASD to hold duplicate data.

Logical Relationships

By establishing a path between two segment types, logical relationships eliminate the need to store duplicate data. To establish a logical relationship, three segment types are always defined:

- A physical parent
- A logical parent
- A logical child

Example: Two databases, one for orders that a customer has placed and one for items that can be ordered, are called ORDER and ITEM. The ORDER database contains information about customers, orders, and delivery. The ITEM database contains information about inventory.

If an application program needs data from both databases, this can be done by defining a logical relationship between the two databases. As shown in Figure 79, a path can be established between the ORDER and ITEM databases using a segment type, called a logical child segment, that points into the ITEM database. Figure 79 is a simple implementation of a logical relationship. In this case, ORDER is the physical parent of ORDITEM. ORDITEM is the physical child of ORDER and the logical child of ITEM.

In a logical relationship, there is a logical parent segment type and it is the segment type pointed to by the logical child. In this example, ITEM is the logical parent of ORDITEM. ORDITEM establishes the path or connection between the two segment types. If an application program now enters the ORDER database, it can access data in the ITEM database by following the pointer in the logical child segment from the ORDER to the ITEM database.

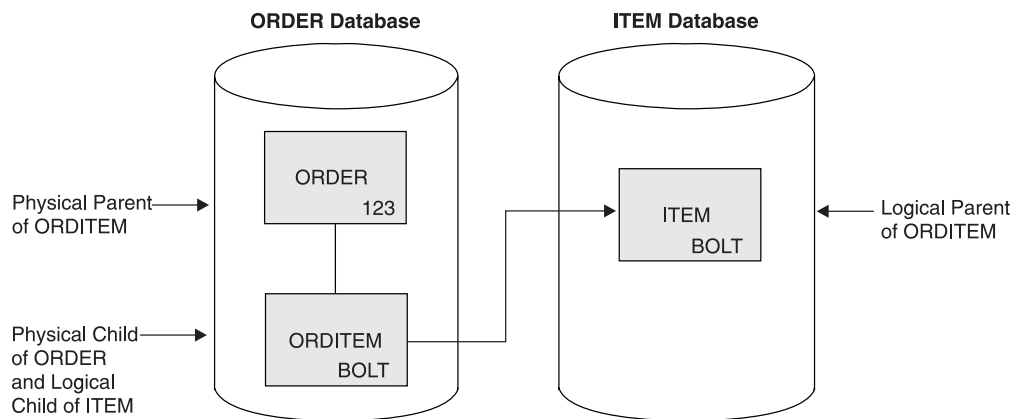


Figure 79. A Simple Logical Relationship

The physical parent and logical parent are the two segment types between which the path is established. The logical child is the segment type that establishes the path. The path established by the logical child is created using pointers.

Logical Relationship Types

Three types of logical relationships are discussed in this topic:

- Unidirectional logical relationships
- Bidirectional physically paired logical relationships
- Bidirectional virtually paired logical relationships

Unidirectional Logical Relationships

A unidirectional relationship links two segment types, a logical child and its logical parent, in one direction. A one-way path is established using a pointer in the logical child. Figure 80 shows a unidirectional relationship that has been established between the ORDER and ITEM databases. A unidirectional relationship can be established between two segment types in the same or different databases. Typically, however, a unidirectional relationship is created between two segment types in different databases. In the figure, the logical relationship can be used to cross from the ORDER to the ITEM database. It cannot be used to cross from the ITEM to the ORDER database, because the ITEM segment does not point to the ORDER database.

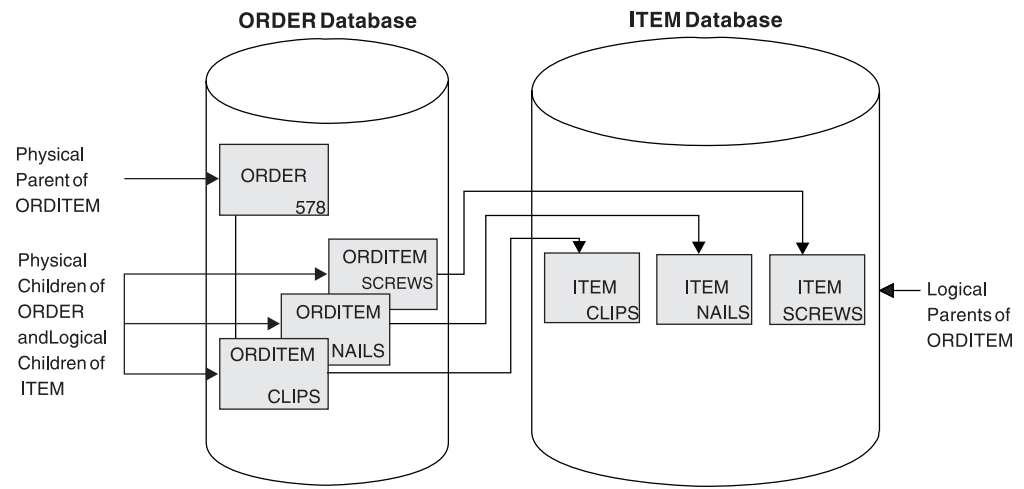


Figure 80. Unidirectional Logical Relationship

It is possible to establish two unidirectional relationships, as shown in Figure 81 on page 160. Then either physical database can be entered and the logical child in either can be used to cross to the other physical database. However, IMS treats each unidirectional relationship as a one-way path. It does not maintain data on both paths. If data in one database is inserted, deleted, or replaced, the corresponding data in the other database is not updated. If, for example, DL/I replaces **ORDITEM-SCREWS** under **ORDER-578**, **ITEMORD-578** under **ITEM-SCREWS** is not replaced. This maintenance problem does not exist in both bidirectional physically paired-logical and bidirectional virtually paired-logical relationships. Both relationship types are discussed next. IMS allows either physical database to be entered and updated and automatically updates the corresponding data in the other database.

Logical Relationships

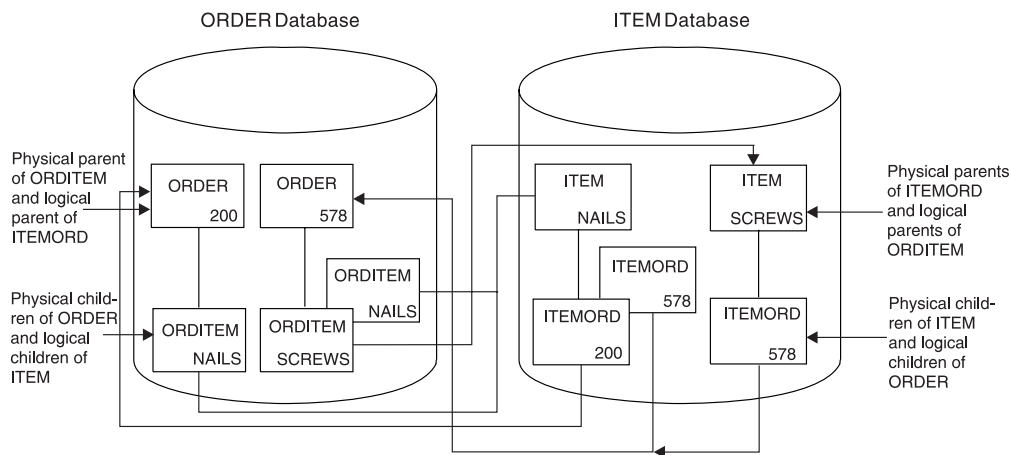


Figure 81. Two Unidirectional Logical Relationships

Bidirectional Physically Paired Logical Relationship

A bidirectional physically paired relationship links two segment types, a logical child and its logical parent, in two directions. A two-way path is established using pointers in the logical child segments. Figure 82 shows a bidirectional physically paired logical relationship that has been established between the ORDER and ITEM databases.

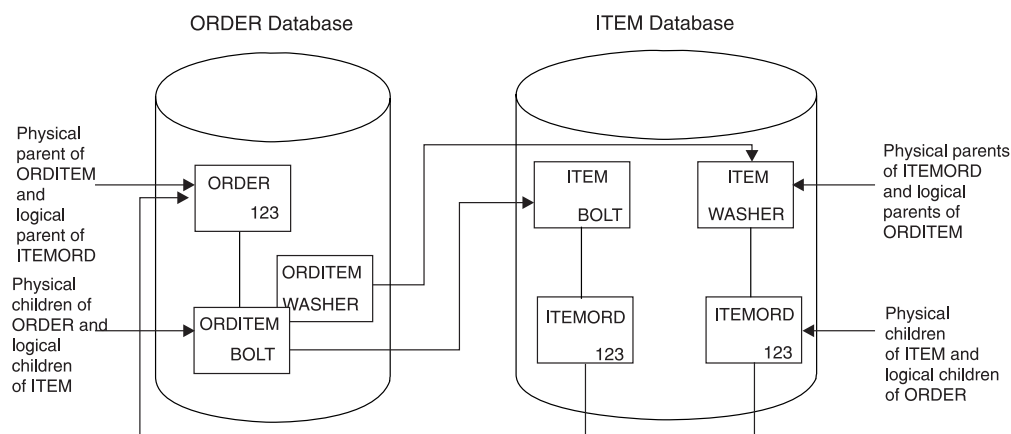


Figure 82. Bidirectional Physically Paired Logical Relationship

Like the other types of logical relationships, a physically paired relationship can be established between two segment types in the same or different databases. The relationship shown in Figure 82 allows either the ORDER or the ITEM database to be entered. When either database is entered, a path exists using the logical child to cross from one database to the other.

In a physically paired relationship, a logical child is stored in both databases. However, if the logical child has dependents, they are only stored in one database. For example, IMS maintains data in both paths in physically paired relationships. In Figure 82 if ORDER 123 is deleted from the ORDER database, IMS deletes from the ITEM database all ITEMORD segments that point to the ORDER 123 segment. If data is changed in a logical child segment, IMS changes the data in its paired

logical child segment. Or if a logical child segment is inserted into one database, IMS inserts a paired logical child segment into the other database.

With physical pairing, the logical child is duplicate data, so there is some increase in storage requirements. In addition, there is some extra maintenance required because IMS maintains data on two paths. In the next type of logical relationship examined, this extra space and maintenance do not exist; however, IMS still allows you to enter either database. IMS also performs the maintenance for you.

Bidirectional Virtually Paired Logical Relationship

A bidirectional virtually paired relationship is like a bidirectional physically paired relationship in that:

- It links two segment types, a logical child and its logical parent, in two directions, establishing a two-way path.
- It can be established between two segment types in the same or different databases.

Figure 83 shows a bidirectional virtually paired relationship between the ORDER and ITEM databases. Note that although there is a two-way path, a logical child segment exists only in the ORDER database. Going from the ORDER to the ITEM database, IMS uses the pointer in the logical child segment. Going from the ITEM to the ORDER database, IMS uses the pointer in the logical parent, as well as the pointer in the logical child segment.

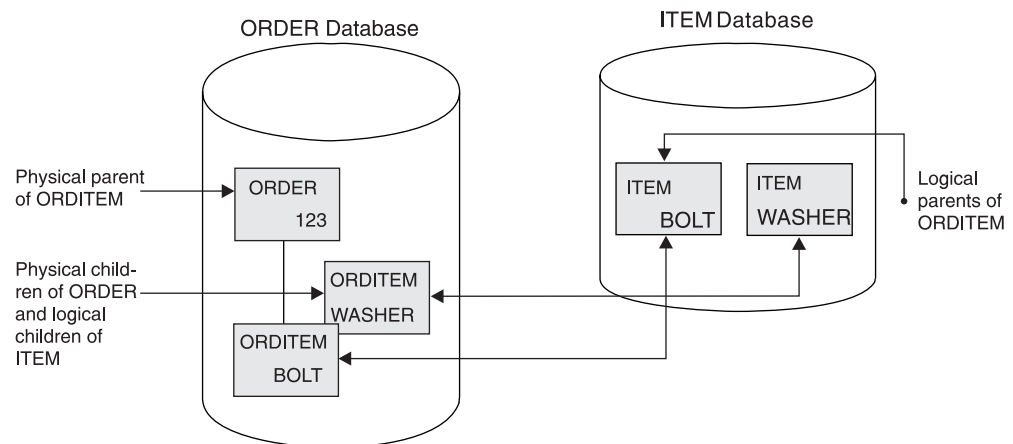


Figure 83. Bidirectionally Virtually Paired Logical Relationship

To define a virtually paired relationship, two logical child segment types are defined in the physical databases involved in the logical relationship. Only one logical child is actually placed in storage. The logical child defined and put in storage is called the *real logical child*. The logical child defined but not put in storage is called the *virtual logical child*.

IMS maintains data in both paths in a virtually paired relationship. However, because there is only one logical child segment, maintenance is simpler than it is in a physically paired relationship. When, for instance, a new ORDER segment is inserted, only one logical child segment has to be inserted. For a replace, the data only has to be changed in one segment. For a delete, the logical child segment is deleted from both paths.

Logical Relationships

Note the trade-off between physical and virtual pairing. With virtual pairing, there is no duplicate logical child and maintenance of paired logical children. However, virtual pairing requires the use and maintenance of additional pointers, called logical twin pointers.

Logical Relationship Pointer Types

In all logical relationships the logical child establishes a path between two segment types. The path is established by use of pointers. The following topics look at pointing in logical relationships and the four types of pointers that you can specify for logical relationships:

- “Logical Parent Pointer”
- “Logical Child Pointer” on page 164
- “Physical Parent Pointer” on page 165
- “Logical Twin Pointer” on page 166

For HALDBs, consider the following:

- Logical relationships are not allowed between HALDBs and non-HALDBs.
- Direct pointers and indirect pointers are used. See “Indirect Pointers” on page 167.
- Unidirectional relationships and bidirectional, physically paired relationships are supported for HALDBs.
- Physical parent pointers are always present in PHDAM and PHIDAM segments.

Logical Parent Pointer

The pointer from the logical child to its logical parent is called a logical parent (LP) pointer. This pointer must be a symbolic pointer when it is pointing into a HISAM database. It can be either a *direct* or a *symbolic* pointer when it is pointing into an HDAM or a HIDAM database. PHDAM or PHIDAM databases require direct pointers.

A direct pointer consists of the direct address of the segment being pointed to, and it can only be used to point into a database where a segment, once stored, is not moved. This means the logical parent segment must be in an HD (HDAM, PHDAM, HIDAM, or PHIDAM) database, since the logical child points to the logical parent segment. The logical child segment, which contains the pointer, can be in a HISAM or an HD database except in the case of HALDB. In the HALDB case, the logical child segment must be in an HD (PHDAM or PHIDAM) database. A direct LP pointer is stored in the logical child's prefix, along with any other pointers, and is four bytes long. Figure 84 on page 163 shows the use of a direct LP pointer. In a HISAM database, pointers are not required between segments because they are stored physically adjacent to each other in hierarchic sequence. Therefore, the only time direct pointers will exist in a HISAM database is when there is a logical relationship using direct pointers pointing into an HD database.

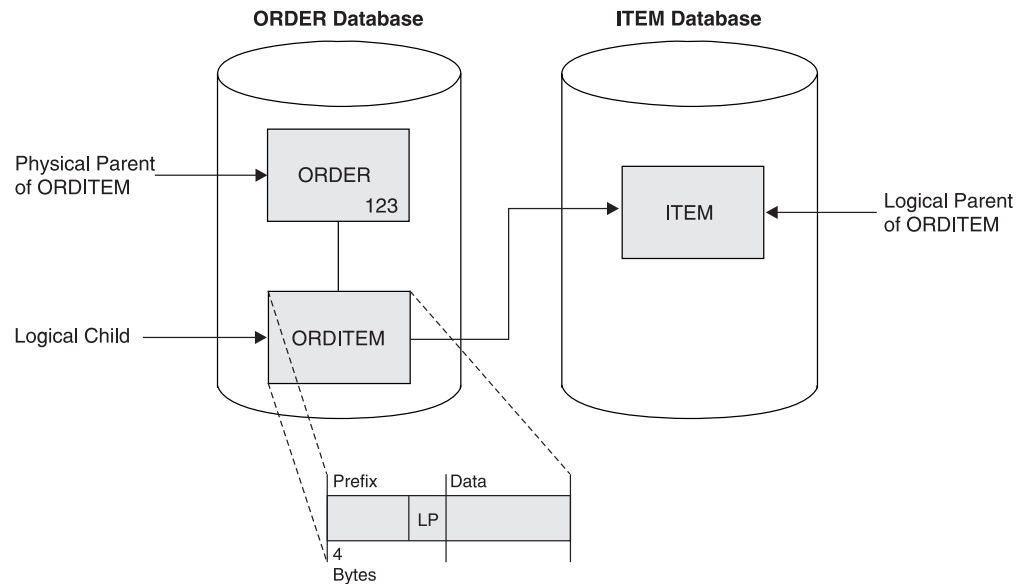


Figure 84. Direct Logical Parent (LP) Pointer

In Figure 84, the direct LP pointer points from the logical child ORDITEM to the logical parent ITEM. Because it is direct, the LP pointer can only point to an HD database. However, the LP pointer can “exist” in a HISAM or an HD database. The LP pointer is in the prefix of the logical child and consists of the 4-byte direct address of the logical parent.

A symbolic LP pointer, which consists of the logical parent’s concatenated key (LPCK), can be used to point into a HISAM or HD database. Figure 85 on page 164 illustrates how to use a symbolic LP pointer. The logical child ORDITEM points to the ITEM segment for BOLT. BOLT is therefore stored in ORDITEM in the LPCK. A symbolic LP pointer is stored in the first part of the data portion in the logical child segment.

Note: The LPCK part of the logical child segment is considered non-replaceable and is not checked to see whether the I/O area is changed. When the LPCK is virtual, checking for a change in the I/O area causes a performance problem. Changing the LPCK in the I/O area does not cause the REPL call to fail. However, the LPCK is not changed in the logical child segment.

With symbolic pointers, if the database the logical parent is in is HISAM or HIDAM, IMS uses the symbolic pointer to access the index to find the correct logical parent segment. If the database containing the logical parent is HDAM, the symbolic pointer must be changed by the randomizing module into a block and RAP address to find the logical parent segment. IMS accesses a logical parent faster when direct pointing is used.

Although the figures show the LP pointer in a unidirectional relationship, it works exactly the same way in all three types of logical relationships.

Figure 85 on page 164 shows an example of a symbolic logical parent pointer.

Logical Relationships

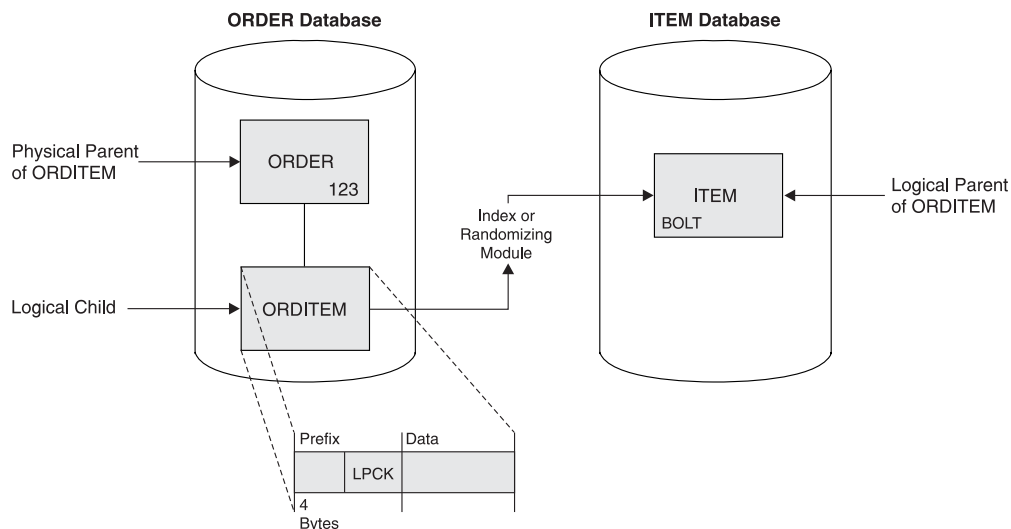


Figure 85. Symbolic Logical Parent (LP) Pointer

In Figure 85, the *symbolic* LP pointer points from the logical child ORDITEM to the logical parent ITEM. With symbolic pointing, the ORDER and ITEM databases can be either HISAM or HD. The LPCK, which is in the first part of the data portion of the logical child, functions as a pointer from the logical child to the logical parent, and is the pointer used in the logical child.

Logical Child Pointer

Logical child pointers are only used in logical relationships with virtual pairing. When virtual pairing is used, there is only one logical child on DASD, called the real logical child. This logical child has an LP pointer. The LP pointer can be symbolic or direct. In the ORDER and ITEM databases you have seen, the LP pointer allows you to go from the database containing the logical child to the database containing the logical parent. To enter either database and cross to the other with virtual pairing, you use a logical child pointer in the logical parent. Two types of logical child pointers can be used:

- Logical child first (LCF) pointers, or
- The *combination* of logical child first (LCF) and logical child last (LCL) pointers

The LCF pointer points from a logical parent to the first occurrence of each of its logical child types. The LCL pointer points to the last occurrence of the logical child segment type for which it is specified. A LCL pointer can only be specified in conjunction with a LCF pointer. Figure 86 on page 165 shows the use of the LCF pointer. These pointers allow you to cross from the ITEM database to the logical child ORDITEM in the ORDER database. However, although you are able to cross databases using the logical child pointer, you have only gone from ITEM to the logical child ORDITEM. To go to the ORDER segment, use the physical parent pointer explained in “Physical Parent Pointer” on page 165.

LCF and LCL pointers are direct pointers. They contain the 4-byte direct address of the segment to which they point. This means the logical child segment, the segment being pointed to, must be in an HD database. The logical parent can be in a HISAM or HD database. If the logical parent is in a HISAM database, the logical child segment must point to it using a symbolic pointer. LCF and LCL pointers are stored in the logical parent's prefix, along with any other pointers. Figure 86 shows a LCF pointer.

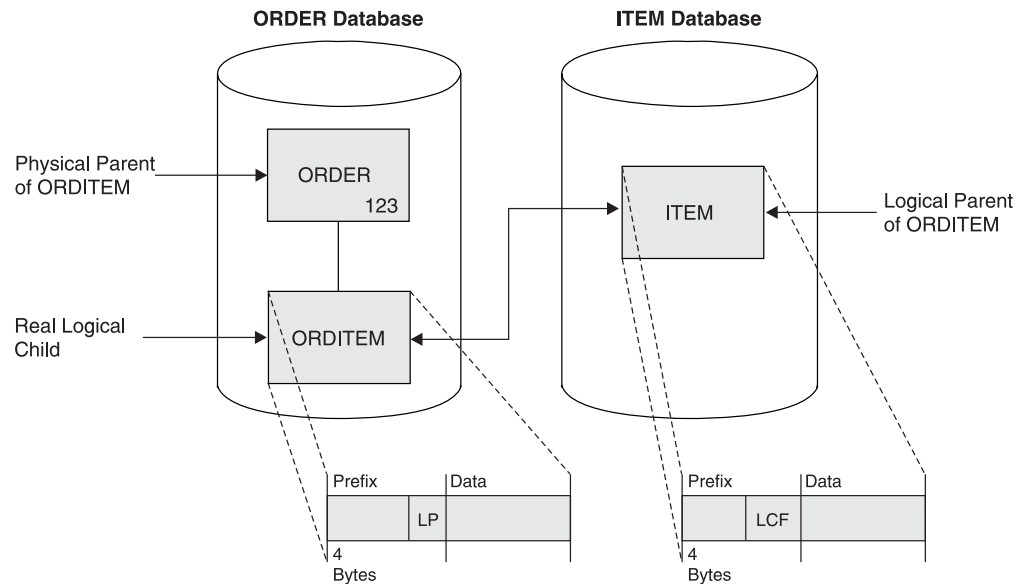


Figure 86. Logical Child First (LCF) Pointer (Used in Virtual Pairing Only)

In Figure 86, the LCF pointer points from the logical parent ITEM to the logical child ORDITEM. Because it is a direct pointer, it can only point to an HD database, although, it can exist in a HISAM or an HD database. The LCF pointer is in the prefix of the logical parent and consists of the 4-byte RBA of the logical child.

Physical Parent Pointer

Physical parent (PP) pointers point from a segment to its physical parent. They are generated automatically by IMS for all HD databases involved in logical relationships. PP pointers are put in the prefix of all logical child and logical parent segments. They are also put in the prefix of all segments on which a logical child or logical parent segment is dependent in its physical database. This creates a path from a logical child or its logical parent back up to the root segment on which it is dependent. Because all segments on which a logical child or logical parent is dependent are chained together with PP pointers to a root, access to these segments is possible in reverse of the usual order.

In Figure 86, you saw that you could cross from the ITEM to the ORDER database when virtual pairing was used, and this was done using logical child pointers. However, the logical child pointer only got you from ITEM to the logical child ORDITEM. Figure 87 on page 166 shows how to get to ORDER. The PP pointer in ORDITEM points to its physical parent ORDER. If ORDER and ITEM are in an HD database but are not root segments, they (and all other segments in the path of the root) would also contain PP pointers to their physical parents.

PP pointers are direct pointers. They contain the 4-byte direct address of the segment to which they point. PP pointers are stored in a logical child or logical parent's prefix, along with any other pointers.

Logical Relationships

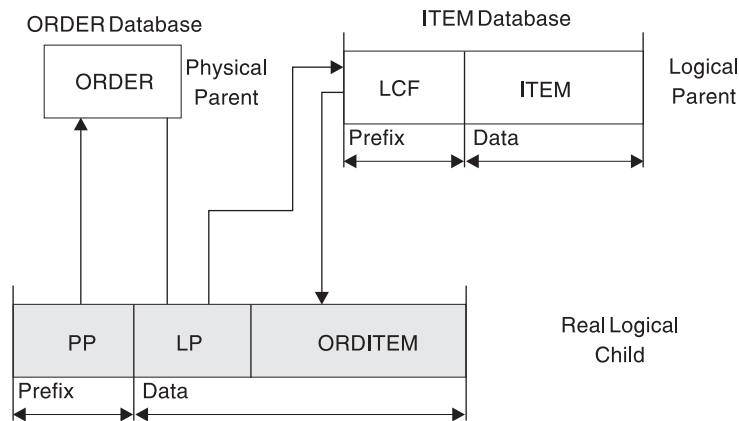


Figure 87. Physical Parent (PP) Pointer

In Figure 87, the PP pointer points from the logical child ORDITEM to its physical parent ORDER. It is generated automatically by IMS for all logical child and logical parent segments in HD databases. In addition, it is in the prefix of the segment that contains it and consists of the 4-byte direct address of its physical parent. PP pointers are generated in all segments from the logical child or logical parent back up to the root.

Logical Twin Pointer

Logical twin pointers are used only in logical relationships with virtual pairing. Logical twins are multiple logical child segments that point to the same occurrence of a logical parent. Two types of logical twin pointers can be used:

- Logical twin forward (LTF) pointers, or
- The combination of logical twin forward (LTF) and logical twin backward (LTB) pointers

An LTF pointer points from a specific logical twin to the logical twin stored after it. An LTB pointer can only be specified in conjunction with an LTF pointer. When specified, an LTB points from a given logical twin to the logical twin stored before it. Logical twin pointers work in a similar way to the physical twin pointers used in HD databases. As with physical twin backward pointers, LTB pointers improve performance on delete operations. They do this when the delete that causes DASD space release is a delete from the physical access path. Similarly, PTB pointers improve performance when the delete that causes DASD space release is a delete from the logical access path.

Figure 88 on page 167 shows use of the LTF pointer. In this example, ORDER 123 has two items: bolt and washer. The ITEMORD segments beneath the two ITEM segments use LTF pointers. If the ORDER database is entered, it can be crossed to the ITEMORD segment for bolts in the ITEM database. Then, to retrieve all items for ORDER 123, the LTF pointers in the ITEMORD segment can be followed. In Figure 88 only one other ITEMORD segment exists, and it is for washers. The LTF pointer in this segment, because it is the last twin in the chain, contains zeros.

LTB pointers on dependent segments improve performance when deleting a real logical child in a virtually paired logical relationship. This improvement occurs when the delete is along the physical path.

LTF and LTB pointers are direct pointers. They contain the 4-byte direct address of the segment to which they point. This means LTF and LTB pointers can only exist in HD databases. Figure 88 shows a LTF pointer.

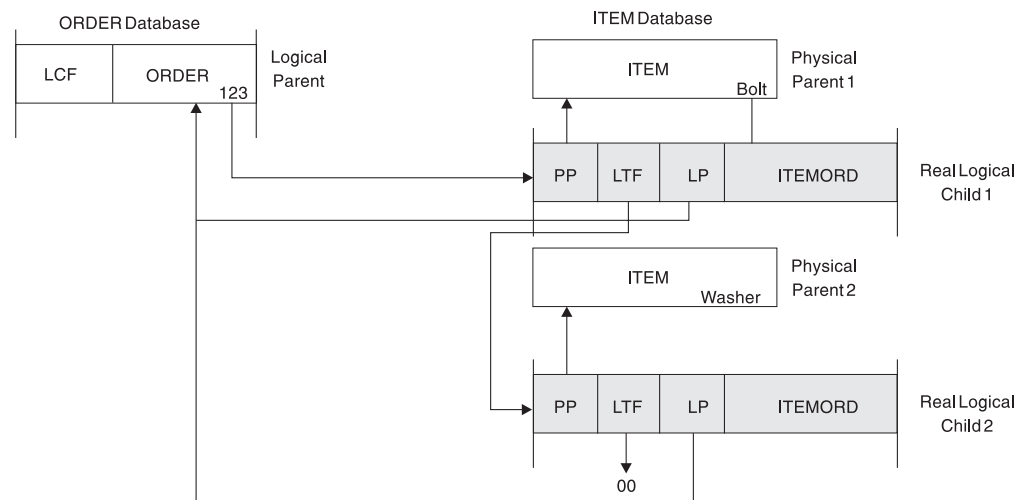


Figure 88. Logical Twin Forward (LTF) Pointer (Used in Virtual Pairing Only)

In Figure 88, the LTF pointer points from a specific logical twin to the logical twin stored after it. In this example, it points from the ITEMORD segment for bolts to the ITEMORD segment for washers. Because it is a direct pointer, the LTF pointer can only point to an HD database. The LTF pointer is in the prefix of a logical child segment and consists of the 4-byte RBA of the logical twin stored after it.

Indirect Pointers

HALDBs (PHDAM, PHIDAM, and PSINDEX databases) use direct and indirect pointers for pointing from one database record to another database record. Figure 89 shows how indirect pointers are used.

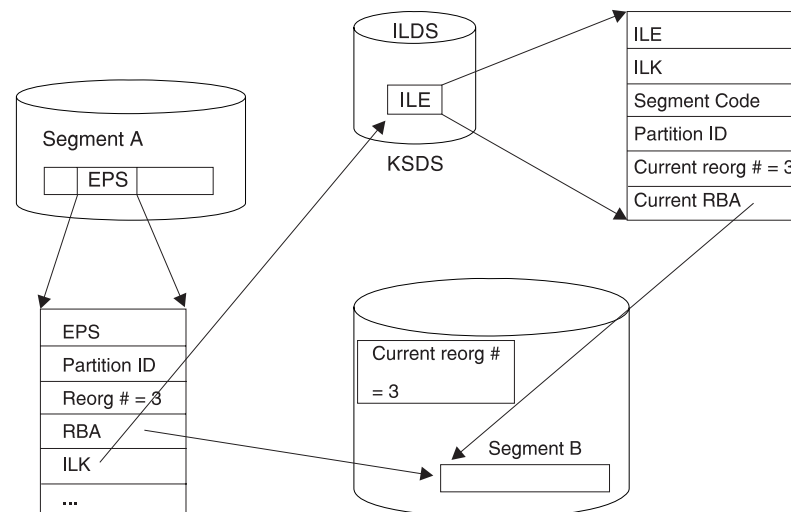


Figure 89. Self-healing Pointers

The use of indirect pointers prevents the problem of misdirected pointers that would otherwise occur when a database is reorganized.

Logical Relationships

The repository for the indirect pointers is the indirect list data set. The misdirected pointers after reorganization are self-healing using indirect pointers.

Related Reading: For a complete discussion of the self-healing pointer process, see “The HALDB Self-Healing Pointer Process” on page 388.

Paths in Logical Relationships

The relationship between physical parent and logical child in a physical database and the LP pointer in each logical child creates a physical parent to logical parent path. To define use of the path, the logical child and logical parent are defined as a concatenated segment type that is a physical child of the physical parent, as shown in Figure 90. Definition of the path and the concatenated segment type is done in what is called a logical database. The logical database is examined in “Specifying Logical Relationships in the Logical DBD” on page 182 and elsewhere in this chapter.




Figure 90. Defining a Physical Parent to Logical Parent Path in a Logical Database

In addition, when LC pointers are used in the logical parent and logical twin and PP pointers are used in the logical child, a logical parent to physical parent path is created. To define use of the path, the logical child and physical parent are defined as one concatenated segment type that is a physical child of the logical parent, as shown in Figure 91. Again, definition of the path is done in a logical database.

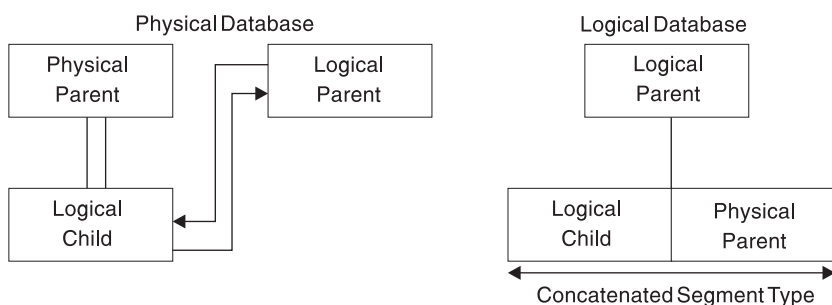


Figure 91. Defining a Logical Parent to Physical Parent Path in a Logical Database

When use of a physical parent to logical parent path is defined, the physical parent is the parent of the concatenated segment type. When an application program retrieves an occurrence of the concatenated segment type from a physical parent, the logical child and its logical parent are concatenated and presented to the application program as one segment. When use of a logical parent to physical parent path is defined, the logical parent is the parent of the concatenated segment type. When an application program retrieves an occurrence of the concatenated

segment type from a logical parent, an occurrence of the logical child and its physical parent are concatenated and presented to the application program as one segment.

In both cases, the physical parent or logical parent segment included in the concatenated segment is called the *destination parent*. For a physical parent to logical parent path, the logical parent is the destination parent in the concatenated segment. For a logical parent to physical parent path, the physical parent is the destination parent in the concatenated segment.

The Logical Child Segment

When defining a logical child in its physical database, the length specified for it must be large enough to contain the concatenated key of the logical parent. Any length greater than that can be used for intersection data.

Related Reading For information about intersection data, see “Intersection Data” on page 170.

To identify which logical parent is pointed to by a logical child, the concatenated key of the logical parent must be present. Each logical child segment must be present in the application program’s I/O area when the logical child is initially presented for loading into the database. However, if the logical parent is in an HD database, its concatenated key might not be written to storage when the logical child is loaded. If the logical parent is in a HISAM database, a logical child in storage must contain the concatenated key of its logical parent.

For logical child segments, you can define a special operand on the PARENT= parameter of the SEGM statement. This operand determines whether a symbolic pointer to the logical parent is stored as part of the logical child segment on the storage device. If PHYSICAL is specified, the concatenated key of the logical parent is stored with each logical child segment. If VIRTUAL is specified, only the intersection data portion of each logical child segment is stored.

When a concatenated segment is retrieved through a logical database, it contains the logical child segment, which consists of the concatenated key of the destination parent, followed by any intersection data. In turn, this is followed by data in the destination parent. Figure 92 shows the format of a retrieved concatenated segment in the I/O area. The concatenated key of the destination parent is returned with each concatenated segment to identify which destination parent was retrieved. IMS gets the concatenated key from the logical child in the concatenated segment or by constructing the concatenated key. If the destination parent is the logical parent and its concatenated key has not been stored with the logical child, IMS constructs the concatenated key and presents it to the application program. If the destination parent is the physical parent, IMS must always construct its concatenated key.




Figure 92. Format of a Concatenated Segment Returned to User I/O Area

Segment Prefix Information for Logical Relationships

There are two things that you should be aware of regarding the prefix of a segment involved in a logical relationship. First, IMS places pointers in the prefix in a specific sequence and, second, IMS places a counter in the prefix for logical parents that do not have logical child pointers.

Sequence of Pointers in a Segment's Prefix

When a segment contains more than one type of pointer and is involved in a logical relationship, pointers are put in the segment's prefix in the following sequence:

1. HF
2. HB
3. PP
4. LTF
5. LTB
6. LP

Or:

1. TF
2. TB
3. PP
4. LTF
5. LTB
6. LP
7. PCF
8. PCL

Or:

1. TF
2. TB
3. PP
4. PCF
5. PCL
6. EPS

Multiple PCF and PCL pointers can exist in a segment type; however, more than one of the other types of pointers can not.

Counter Used in Logical Relationships

IMS puts a 4-byte counter in all logical parents that do not have logical child pointers. The counter is stored in the logical parent's prefix and contains a count of the number of logical children pointing to this logical parent. The counter is maintained by IMS and is used to handle delete operations properly. If the count is greater than zero, the logical parent cannot be deleted from the database because there are still logical children pointing to it.

Intersection Data

When two segments are logically related, data can exist that is unique to only that relationship. In Figure 93 on page 171, for example, one of the items ordered in ORDER 123 is 5000 bolts. The quantity 5000 is specific to this order (ORDER 123) and this item (bolts). It does not belong to either the order or item on its own.

Similarly, in ORDER 123, 6000 washers are ordered. Again, this data is concerned only with that particular order and item combination.

This type of data is called intersection data, since it has meaning only for the specific logical relationship. The quantity of an item could not be stored in the ORDER 123 segment, because different quantities are ordered for each item in ORDER 123. Nor could it be stored in the ITEM segment, because for each item there can be several orders, each requesting a different quantity. Because the logical child segment links the ORDER and ITEM segments together, data that is unique to the relationship between the two segments can be stored in the logical child.

The two types of intersection data are: fixed intersection data (FID) and variable intersection data (VID).

Fixed Intersection Data

Data stored in the logical child is called fixed intersection data (FID). When symbolic pointing is used, it is stored in the data part of the logical child after the LPCK. When direct pointing is used, it is the only data in the logical child segment. Because symbolic pointing is used in Figure 93, BOLT and WASHER are the LPCK, and the 5000 and 6000 are the FID. The FID can consist of several fields, all of them residing in the logical child segment.

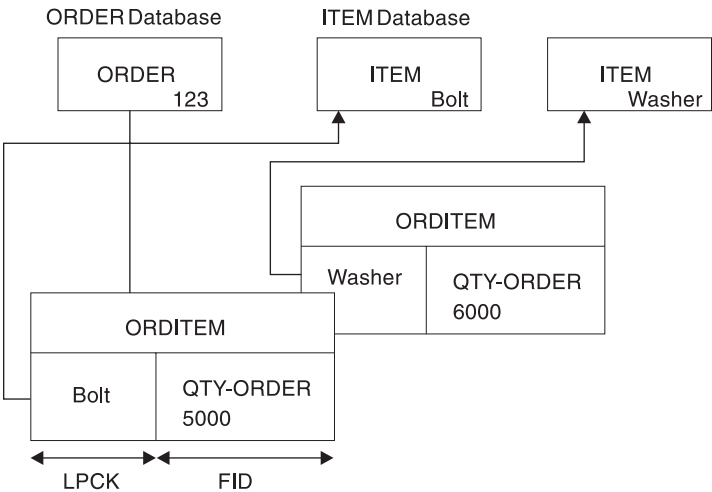


Figure 93. Fixed Intersection Data

Variable Intersection Data

VID is used when you have data that is unique to a relationship, but several occurrences of it exist. For example, suppose you cannot supply in one shipment the total quantity of an item required for an order. You need to store delivery data showing the quantity delivered on a specified date. The delivery date is not dependent on either the order or item alone. It is dependent on a specific order-item combination. Therefore, it is stored as a dependent of the logical child segment. The data in this dependent of the logical child is called variable intersection data. For each logical child occurrence, there can be as many occurrences of dependent segments containing intersection data as you need.

|
|
|

Figure 94 on page 172 shows variable intersection data. In the ORDER 123 segment for the item BOLT, 3000 were delivered on March 2 and 1000 were delivered on April 2. Because of this, two occurrences of the DELIVERY segment

Logical Relationships

exist. Multiple segment types can contain intersection data for a single logical child segment. In addition to the DELIVERY segment shown in the figure, note the SCHEDULE segment type. This segment type shows the planned shipping date and the number of items to be shipped. Segment types containing VID can all exist at the same level in the hierarchy as shown in the figure, or they can be dependents of each other.

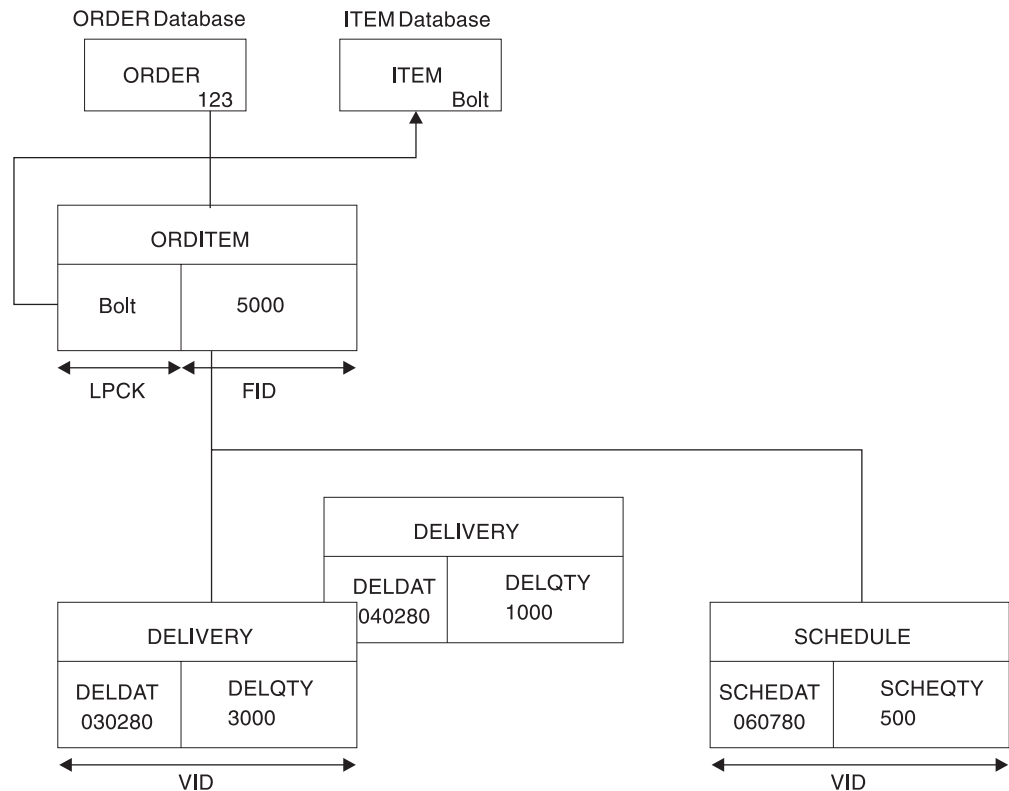


Figure 94. Variable Intersection Data

FID, VID, and Physical Pairing

In the previous figures, intersection data has been stored in a unidirectional logical relationship. It works exactly the same way in the two bidirectional logical relationships. However, when physical pairing is used, VID can only be stored on one side of the relationship. It does not matter on which side it is stored. An application program can access it using either the ORDER or ITEM database. FID, on the other hand, must be stored on both sides of the relationship when physical pairing is used. IMS automatically maintains the FID on both sides of the relationship when it is changed on one side. However, extra time is required for maintenance, and extra space is required on DASD for FID in a physically paired relationship.

Recursive Structures: Same Database Logical Relationships

Logical relationships can be established between segments in two or more physical databases. Logical relationships can also be established between segments in the same database. The logical data structure that results is called a *recursive structure*.

Most often, recursive structures are defined in manufacturing for bill-of-materials type applications. Suppose, for example, a company manufactures bicycles. The

first model the manufacturer makes is Model 1, which is a boy's bicycle. Table 16 lists the parts needed to manufacture this bicycle and the number of each part needed to manufacture one Model 1 bicycle.

Table 16. Parts List for the Model 1 Bicycle Example

| Part | Number Needed |
|-----------------------|---------------|
| 21-inch boy's frame | 1 |
| Handlebar | 1 |
| Seat | 1 |
| Chain | 1 |
| Front fender | 1 |
| Rear fender | 1 |
| Pedal | 2 |
| Crank | 1 |
| Front sprocket | 1 |
| 26-inch tube and tire | 2 |
| 26-inch rim | 2 |
| 26-inch spoke | 72 |
| Front hub | 1 |
| Housing | 1 |
| Break | 1 |
| Rear sprocket | 1 |

In manufacturing, it is necessary to know the steps that must be executed to manufacture the end product. For each step, the parts needed must be available and any subassemblies used in a step must have been assembled in previous steps. Figure 95 on page 174 shows the steps required to manufacture the Model 1 bicycle. A housing, brake, and rear sprocket are needed to make the rear hub assembly in step 2. Only then can the part of step 3 that involves building the rear wheel assembly be executed. This part of step 3 also requires availability of a 26-inch tire, a rim, and 36 spokes.

The same company manufactures a Model 2 bicycle, which is for girls. The parts and assembly steps for this bicycle are exactly the same, except that the bicycle frame is a girl's frame.

If the manufacturer stored all parts and subassemblies for both models as separate segments in the database, a great deal of duplicate data would exist. Figure 95 on page 174 shows the segments that must be stored just for the Model 1 bicycle. A similar set of segments must be stored for the Model 2 bicycle, except that it has a girl's bicycle frame. As you can see, this leads to duplicate data and the associated maintenance problems. The solution to this problem is to create a recursive structure. Figure 96 on page 175 shows how this is done using the data for the Model 1 bicycle.

Logical Relationships

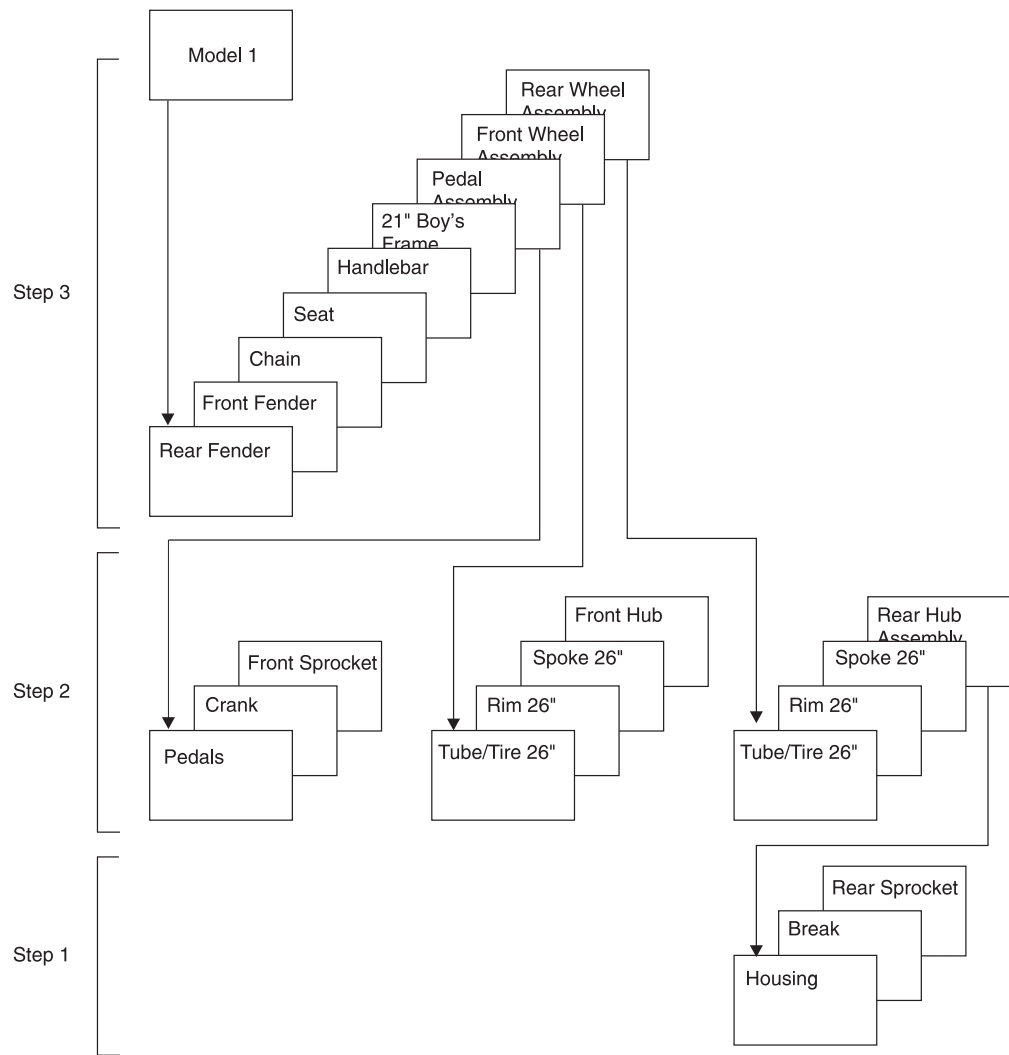


Figure 95. Model 1 Components and Subassemblies

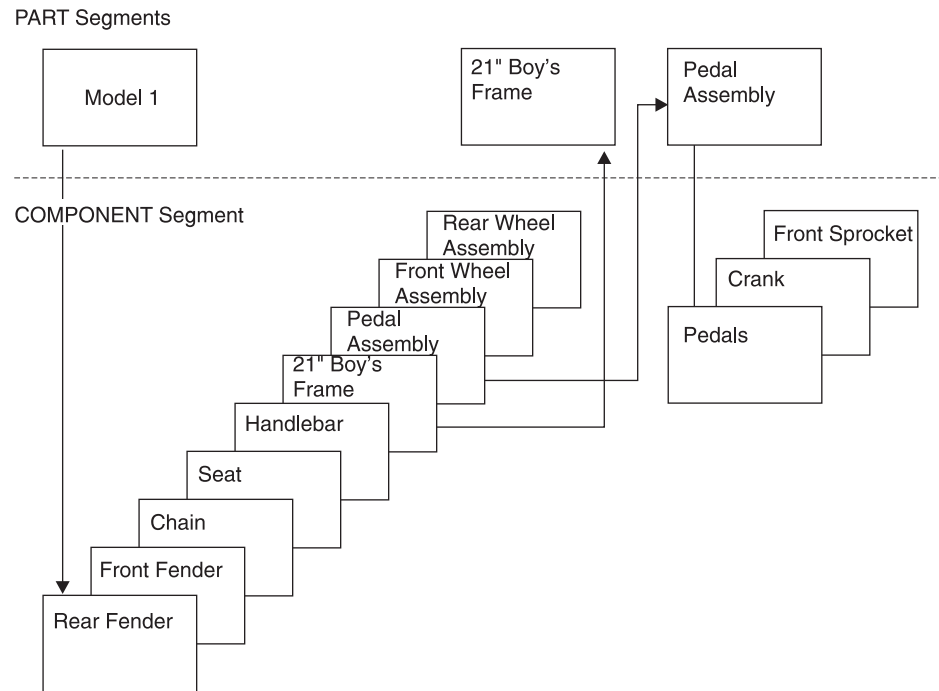


Figure 96. Database Records for the Model 1 Bicycle

In Figure 96, two types of segments exist: PART and COMPONENT segments. A unidirectional logical relationship has been established between them. The PART segment for Model 1 is a root segment. Beneath it are nine occurrences of COMPONENT segments. Each of these is a logical child that points to another PART root segment. (Only two of the pointers are actually shown to keep the figure simple.) However, the other PART root segments show the parts required to build the component.

For example, the pedal assembly component points to the PART root segment for assembling the pedal. Stored beneath this segment are the following parts that must be assembled: one front sprocket, one crank, and two pedals. With this structure, much of the duplicate data otherwise stored for the Model 2 bicycle can be eliminated.

Figure 97 on page 176 shows the segments, in addition to those in Figure 96, that must be stored in the database record for the Model 2 bicycle. The logical children in the figure, except the one for the unique component, a 21" girl's frame, can point to the same PART segments as are shown in Figure 96. A separate PART segment for the pedal assembly, for example, need not exist. The database record for both Model 1 and 2 have the same pedal assembly, and by using the logical child, it can point to the same PART segment for the pedal assembly.

Logical Relationships

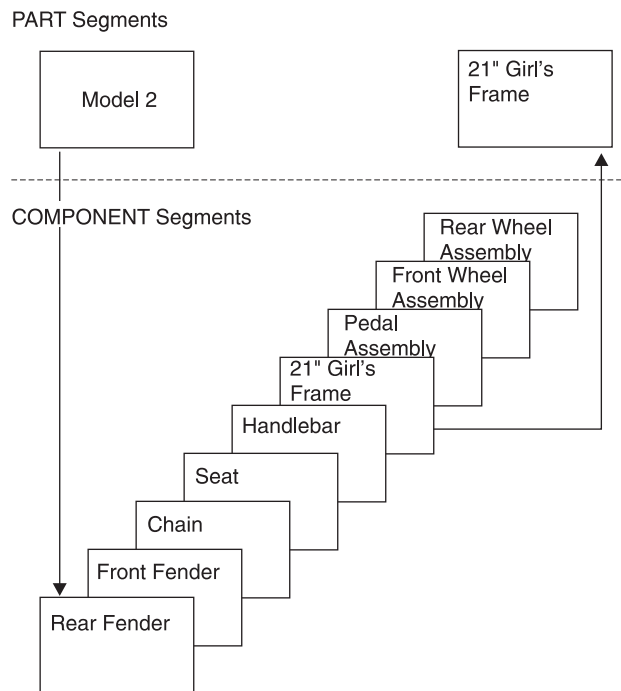


Figure 97. Extra Database Records Required for the Model 2 Bicycle

One thing to note about recursive structures is that the physical parent and the logical parent of the logical child are the same segment type. For example, in Figure 96 on page 175, the PART segment for *Model 1* is the physical parent of the COMPONENT segment for pedal assembly. The PART segment for pedal assembly is the logical parent of the COMPONENT segment for pedal assembly.

Defining Sequence Fields for Logical Relationships

This topic discusses defining the following types of sequence fields:

- “Logical Parent Sequence Fields”
- “Real Logical Children Sequence Fields” on page 177
- “Virtual Logical Children Sequence Fields” on page 177

Logical Parent Sequence Fields

To avoid potential problems in processing databases using logical relationships, unique sequence fields should be defined in all logical parent segments. In all segments a logical parent is dependent on in its physical database. When unique sequence fields are not defined in all segments on the path to and including a logical parent, multiple logical parents in a database can have the same concatenated key. When this happens, problems can arise during and after initial database load when symbolic logical parent pointers in logical child segments are used to establish position on a logical parent segment.

At initial database load time, if logical parents with non-unique concatenated keys exist in a database, the resolution utilities (described in Chapter 15, “Tuning Databases,” on page 347) attach all logical children with the same concatenated key to the first logical parent in the database with that concatenated key.

When inserting or deleting a concatenated segment and position for the logical parent, part of the concatenated segment is determined by the logical parent’s concatenated key. Positioning for the logical parent starts at the root and stops on

the first segment at each level of the logical parent's database that satisfies the key equal condition for that level. If a segment is missing on the path to the logical parent being inserted, a GE status code is returned to the application program when using this method to establish position in the logical parent's database.

Real Logical Children Sequence Fields

If the sequence field of a real logical child consists of any part of the logical parent's concatenated key, PHYSICAL must be specified on the PARENT= parameter in the SEGM statement for the logical child. This will cause the concatenated key of the logical parent to be stored with the logical child segment.

Virtual Logical Children Sequence Fields

As a general rule, a segment can have only one sequence field. However, in the case of virtual pairing, multiple FIELD statements can be used to define a logical sequence field for the virtual logical child.

A sequence field must be specified for a virtual logical child if, when accessing it from its logical parent, you need real logical child segments retrieved in an order determined by data in a field of the virtual logical child as it could be seen in the application program I/O area. This sequence field can include any part of the segment as it appears when viewed from the logical parent (that is, the concatenated key of the real logical child's physical parent followed by any intersection data). Because it can be necessary to describe the sequence field of a logical child as accessed from its logical parent in non-contiguous pieces, multiple FIELD statements with the SEQ parameter present are permitted. Each statement must contain a unique fldname1 parameter.

Control Blocks for Logical Relationships

When a logical relationship is used, you must define the physical databases involved in the relationship to IMS. This is done using a *physical* DBD. In addition, many times you must define the logical structure of IMS since this is the structure the application program perceives. This is done using a *logical* DBD. A logical DBD is needed because the application program's PCB references a DBD, and the physical DBD does not reflect the logical data structure the application program needs to access. Finally, the application program needs a PSB, consisting of one or more PCBs. The PCB that is used when processing with a logical relationship points to the logical DBD when one has been defined. This PCB indicates which segments in the logical database the application program can process. It also indicates what type of processing the application program can perform on each segment.

Figure 98 on page 178 shows the relationship between these three control blocks. It assumes that the logical relationship is established between two physical databases. The following topics explain how the physical and logical DBD are coded when a logical relationship is defined:

- "Specifying Logical Relationships in the Physical DBD" on page 178
- "Specifying Logical Relationships in the Logical DBD" on page 182

Logical Relationships




Figure 98. Relationship of Control Blocks When a Logical Relationship Is Used

Specifying Logical Relationships in the Physical DBD

For each of the databases involved in a logical relationship, you must code a physical DBD. All statements in the physical DBD are coded with the same format used when a logical relationship is not defined, except for the SEGM and LCHILD statements. The SEGM statement, which describes a segment and its length and position in the database hierarchy, is expanded to include the new types of pointers. The LCHILD statement is added to define the logical relationship between the two segment types. Figure 100 on page 179 shows an example of how the physical DBD is coded.

In the SEGM statements of the examples associated with Figure 99 on page 179 and Figure 100 on page 179, only the pointers required with logical relationships are shown. No pointers required for use with HD databases are shown. When actually coding a DBD, you must ask for these pointers in the PTR= parameter. Otherwise, IMS will not generate them once another type of pointer is specified.

Figure 99 shows the layout of segments. Figure 100 on page 179 shows physical DBDs for unidirectional relationships.

| | | | | |
|--------|-------------------------|----|--------|-------|
| ORDER | | | | |
| ORDKEY | DESCRIPTION OF LOCATION | | ORDATE | * |
| Bytes | 1 | 11 | 41 | 47 50 |

| | |
|---------|----------|
| ORDITEM | |
| ITEMNO | ORDITQTY |
| Bytes | 1 9 17 |

| | | | | |
|----------|----------|-------------|----|----|
| DELIVERY | | | | |
| DELDAT | QUANTITY | DESCRIPTION | * | |
| Bytes | 1 7 | 15 | 45 | 50 |

| | | | | |
|----------|----------|-----------------------|-----------------|-------------|
| SCHEDULE | | | | |
| SCHEDAT | QUANTITY | PLANNED SHIPPING DATE | INVENTORY PLACE | DESCRIPTION |
| Bytes | 1 7 | 15 | 21 25 | 50 |

| | | | | |
|---------|-------------|----|------------------|----|
| ITEM | | | | |
| ITEMKEY | DESCRIPTION | | DATE OF CREATION | * |
| Bytes | 1 9 | 49 | 55 | 60 |

Figure 99. Layouts of Segments Used in the Examples

This is the hierarchic structure of the two databases involved in the logical relationship. In this example, we are defining a unidirectional relationship using symbolic pointing. ORDITEM has an LPCK and FID, and DELIVERY and SCHEDULE are VID.

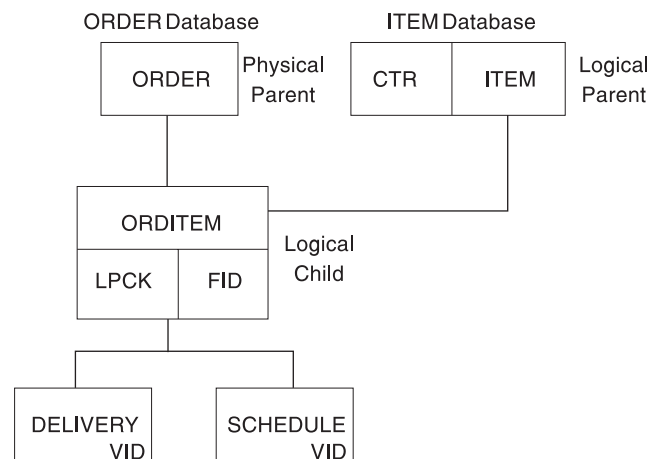


Figure 100. Physical DBDs for Unidirectional Relationship Using Symbolic Pointing

Logical Relationships

The following DBD is for the ORDER database:

```
DBD    NAME=ORDDB
SEGM    NAME=ORDER,BYTES=50,FREQ=28000,PARENT=0
FIELD    NAME=(ORDKEY,SEQ),BYTES=10,START=1,TYPE=C
FIELD    NAME=ORDATE,BYTES=6,START=41,TYPE=C
SEGM    NAME=ORDITEM,BYTES=17,PARENT=((ORDER),(ITEM,P,ITEMDB))
FIELD    NAME=(ITEMNO,SEQ),BYTES=8,START=1,TYPE=C
FIELD    NAME=ORDITQTY,BYTES=9,START=9,TYPE=C,
SEGM    NAME=DELIVERY,BYTES=50,PARENT=ORDITEM
FIELD    NAME=(DELDAT,SEQ),BYTES=6,START=1,TYPE=C
SEGM    NAME=SCHEDULE,BYTES=50,PARENT=ORDITEM
FIELD    NAME=(SCHEDAT,SEQ),BYTES=6,START=1,TYPE=C
DBDGEN
FINISH
END
```

The following DBD is for the ITEM database:

```
DBD    NAME=ITEMDB
SEGM    NAME=ITEM,BYTES=60,FREQ=50000,PARENT=0
FIELD    NAME=(ITEMKEY,SEQ),BYTES=8,START=1,TYPE=C
LCHILD    NAME=(ORDITEM,ORDDB)
DBDGEN
FINISH
END
```

Notes to Figure 100:

In the ORDER database, the DBD coding that differs from normal DBD coding is that for the logical child ORDITEM.

In the SEGM statement for ORDITEM:

1. The BYTES= parameter is 17. The length specified is the length of the LPCK, plus the length of the FID. The LPCK is the key of the ITEM segment, which is 8 bytes long. The length of the FID is 9 bytes.
2. The PARENT= parameter has two parents specified. Two parents are specified because ORDITEM is a logical child and therefore has both a physical and logical parent. The physical parent is ORDER. The logical parent is ITEM, specified after ORDER. Because ITEM exists in a different physical database from ORDITEM, the name of its physical database, ITEMDB, must be specified. Between the segment name ITEM and the database name ITEMDB is the letter P. The letter P stands for physical. The letter P specifies that the LPCK is to be stored on DASD as part of the logical child segment.

In the FIELD statements for ORDITEM:

1. ITEMNO is the sequence field of the ORDITEM segment and is 8 bytes long. ITEMNO is the LPCK. The logical parent is ITEM, and if you look at the FIELD statement for ITEM in the ITEM database, you will see ITEM's sequence field is ITEMKEY, which is 8 bytes long. Because ITEM is a root segment, the LPCK is 8 bytes long.
2. ORDITQTY is the FID and is coded normally.

In the ITEM database, the DBD coding that differs from normal DBD coding is that an LCHILD statement has been added. This statement names the logical child ORDITEM. Because the ORDITEM segment exists in a different physical database from ITEM, the name of its physical database, ORDDb, must be specified.

Specifying Bidirectional Logical Relationships

Figure 100 on page 179 shows the coding for a unidirectional relationship. When defining a bidirectional relationship with physical pairing, you need to include an LCHILD statement under both logical parents. In addition to other pointers, you need to include the PAIRED operand on the POINTER= parameter of the SEGM statements for both logical children.

When defining a bidirectional relationship with virtual pairing, you need to code an LCHILD statement only for the real logical child. On the LCHILD statement, you code POINTER=SNGL or DBLE to get logical child pointers. You code the PAIR= operand to indicate the virtual logical child that is paired with the real logical child. When you define the SEGM statement for the real logical child, the PARENT= parameter identifies both the physical and logical parents. You should specify logical twin pointers (in addition to any other pointers) on the POINTER= parameter. Also, you should define a SEGM statement for the virtual logical child even though it does not exist. On this SEGM statement, you specify PAIRED on the POINTER= parameter. In addition, you specify a SOURCE= parameter. On the SOURCE= parameter, you specify the SEGM name and DBD name of the real logical child. DATA must always be specified when defining SOURCE= on a virtual logical child SEGM statement.

Related Reading: For more information on coding logical relationships, see *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

Checklist of Rules for Defining Logical Relationships in Physical Databases

This topic provides the list of rules that must be followed when defining logical relationships in physical databases. In all cases, references are to segment types, *not* occurrences.

Logical Child Rules:

- A logical child must have a physical and a logical parent.
- A logical child can have only one physical and one logical parent.
- A logical child is defined as a physical child in the physical database of its physical parent.
- A logical child is always a dependent segment in a physical database, and can, therefore, be defined at any level except the first level of a database.
- A logical child in its physical database cannot have a physical child defined at the next lower level in the database that is also a logical child.
- A logical child can have a physical child. However, if a logical child is physically paired with another logical child, only one of the paired segments can have physical children.

Logical Parent Rules:

- A logical parent can be defined at any level in a physical database, including the root level.
- A logical parent can have one or more logical children. Each logical child related to the same logical parent defines a logical relationship.
- A segment in a physical database cannot be defined as both a logical parent and a logical child.
- A logical parent can be defined in the same physical database as its logical child, or in a different physical database.

Rules for Defining Logical Relationships

Physical Parent Rules: A physical parent of a logical child cannot also be a logical child.

Specifying Logical Relationships in the Logical DBD

To identify which segment types are used in a logical data structure, you must code a logical DBD. Figure 101 shows an example of how the logical DBD is coded. The example is a logical DBD for the same physical databases defined in “Specifying Logical Relationships in the Physical DBD” on page 178.

When defining a segment in a logical database, you can specify whether the segment is returned to the program’s I/O area by using the KEY or DATA operand on the SOURCE= parameter of the SEGM statement. DATA returns both the key and data portions of the segment to the I/O area. KEY returns only the key portion, and not the data portion of the segment to the I/O area.

When the SOURCE= parameter is used on the SEGM statement of a concatenated segment, the KEY and DATA parameters control which of the two segments, or both, is put in the I/O area on retrieval calls. In other words, you define the SOURCE= parameter twice for a concatenated segment type, once for the logical child portion and once for the destination parent portion.

Figure 101 illustrates the logical data structure you need to create in the application program. It is implemented with a unidirectional logical relationship using symbolic pointing. The root segment is ORDER from the ORDER database. Dependent on ORDER is ORDIT, the logical child, concatenated with its logical parent ITEM. The application program receives both segments in its I/O area when a single call is issued for ORDIT. DELIVERY and SCHEDULE are VID.

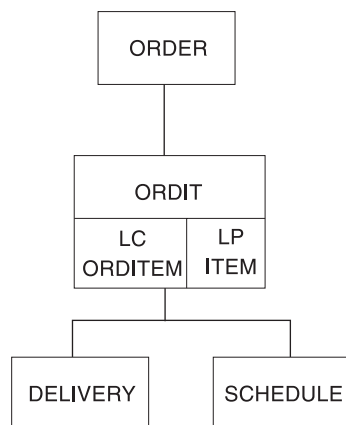


Figure 101. Logical Data Structure for a Unidirectional Relationship Using Symbolic Pointing

The following logical DBD is for the logical data structure shown in Figure 101:

```
DBD NAME=ORDLOG,ACCESS=LOGICAL
DATASET LOGICAL
SEGM NAME=ORDER,SOURCE=((ORDER,DATA,ORDDB))
SEGM NAME=ORDIT,PARENT=ORDER,                                     X
      SOURCE=((ORDITEM,DATA,ORDDB),(ITEM,DATA,ITEMDB))
SEGM NAME=DELIVERY,PARENT=ORDIT,SOURCE=((DELIVERY,DATA,ORDDB))
SEGM NAME=SCHEDULE,PARENT=ORDIT,SOURCE=((SCHEDULE,DATA,ORDDB))
DBDGEN
FINISH
END
```

Notes to Figure 101:

1. The DBD statement has the name of the logical DBD, in this example ORDLOG. As with physical DBDs, this name must be unique and must be the same name as specified in the MBR operand of the DBDGEN procedure. ACCESS=LOGICAL simply says this is a logical DBD.
2. The DATASET statement always says LOGICAL, meaning a logical DBD. No other parameters can be specified on this statement; however, DDNAMEs for data sets are all specified in the DATASET statements in the physical DBDs.
3. The SEGM statements show which segments are to be included in the logical database. The only operands allowed on the SEGM statements for a logical DBD are NAME=, PARENT=, and SOURCE=. All other information about the segment is defined in the physical DBD.
 - The first SEGM statement defines the root segment ORDER.
 The NAME= operand specifies the name used in the PCB to refer to this segment. This name is used by application programmers when coding SSAs. In this example, the segment name is the same as the name used in the physical DBD - ORDER. However, the segment could have a different name from that specified in its physical DBD.
 The SOURCE= operand tells IMS where the data for this segment is to come from. First the name of the segment type appears in its physical database, which is ORDER. DATA says that the data in this segment needs to be put in the application program's I/O area. ORDDDB is the name of the physical database in which the ORDER segment exists.
 No FIELD statements are coded in the logical DBD. IMS picks the statements up from the physical DBD, so when accessing the ORDER segment in this logical data structure, the application program could have SSAs referring to ORDKEY or ORDATE. These fields were defined for the ORDER segments in its physical DBD, as shown in Figure 100 on page 179.
 - The second SEGM statement is for the ORDIT segment. The ORDIT segment is made up of the logical child ORDITEM, concatenated with its logical parent ITEM. As you can see, the SOURCE= operand identifies both the ORDITEM and ITEM segments in their different physical databases.
 - The third and fourth SEGM statements are for the VID DELIVERY and SCHEDULE. These SEGM statements must be placed in the logical DBD in the same relative order they appear in the physical DBD. In the physical DBD, DELIVERY is to the left of SCHEDULE.

Checklist of Rules for Defining Logical Databases

Before the rules for defining logical databases can be described, you need to know the following definitions:

- Crossing a logical relationship
- The first and additional logical relationships crossed

Also, a logical DBD is needed only when an application program needs access to a concatenated segment or needs to cross a logical relationship.

Definition of Crossing a Logical Relationship: A logical relationship is considered crossed when it is used in a logical database to access a segment that is:

- A physical parent of a destination parent in the destination parent's database
- A physical dependent of a destination parent in the destination parent's physical database

Rules for Defining Logical Relationships

If a logical relationship is used in a logical database to access a destination parent only, the logical relationship is not considered crossed.

In Figure 102, DBD1 and DBD2 are two physical databases with a logical relationship defined between them. DBD3 through DBD6 are four logical databases that can be defined from the logical relationship between DBD1 and DBD2. With DBD3, no logical relationship is crossed, because no physical parent or physical dependent of a destination parent is included in DBD3. With DBD4 through DBD6, a logical relationship is crossed in each case, because each contains a physical parent or physical dependent of the destination parent.

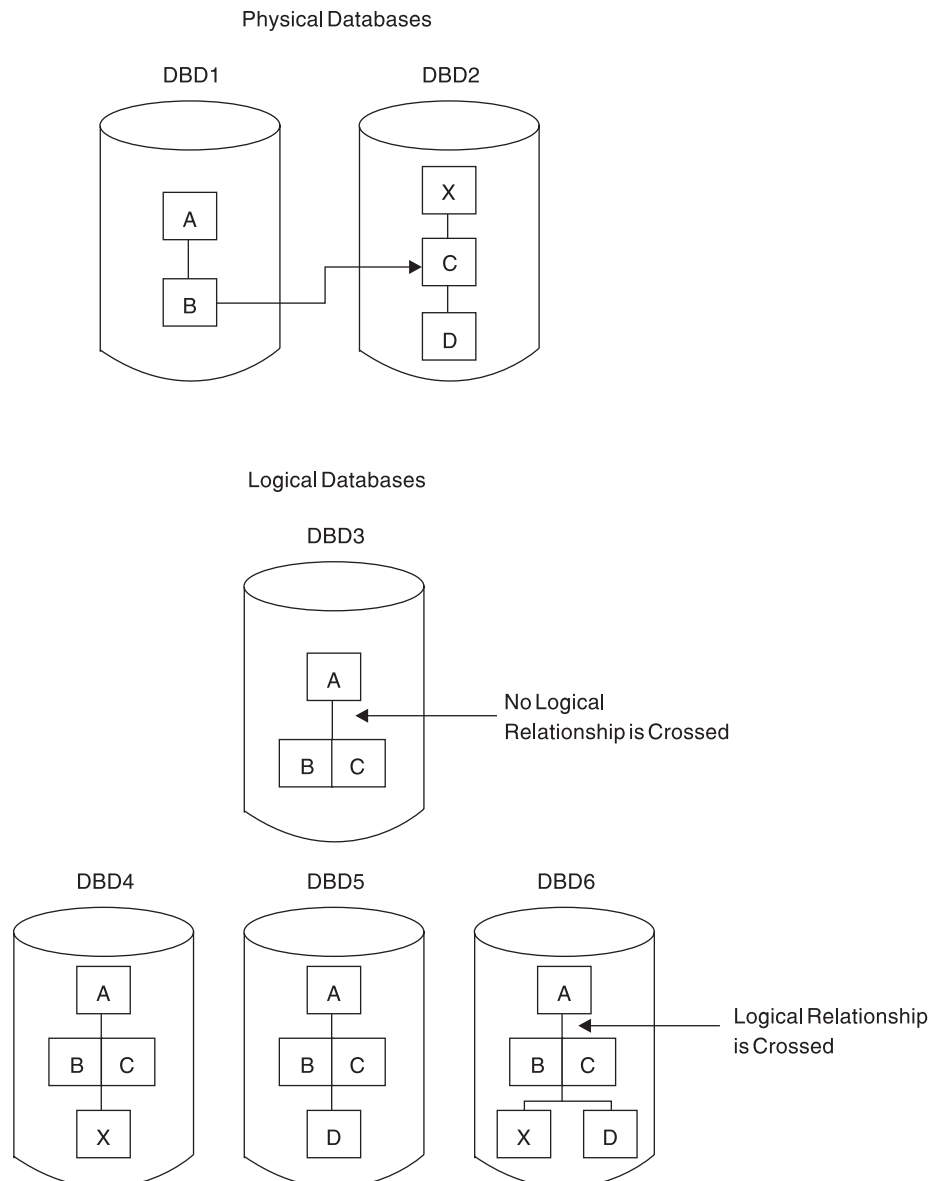


Figure 102. Definition of Crossing a Logical Relationship

Definition of First and Additional Logical Relationships Crossed: More than one logical relationship can be crossed in a hierarchic path in a logical database. Figure 103 on page 185 shows three physical databases (DBD1, DBD2 and DBD3) in which logical relationships have been defined. Also in the figure are two (of many) logical databases (DBD4 and DBD5) that can be defined from the logical

relationships in the physical databases. In DBD4, the two concatenated segments BF and DI allow access to all segments in the hierarchic paths of their destination parents. If either logical relationship or both is crossed, each is considered the first logical relationship crossed. This is because each concatenated segment type is reached by following the physical hierarchy of segment types in DBD1.

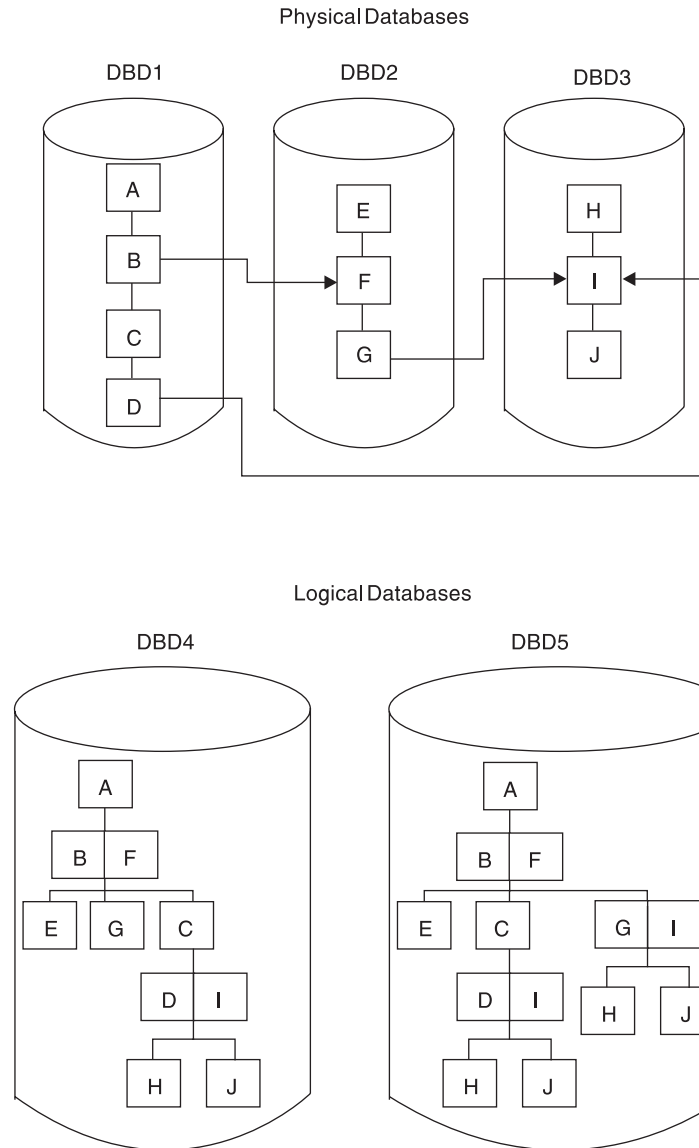


Figure 103. The First Logical Relationship Crossed in a Hierarchic Path of a Logical Database

In DBD5 in Figure 103, an additional concatenated segment type GI, is defined that was not included in DBD4. GI allows access to segments in the hierarchic path of the destination parent if crossed. When the logical relationship made possible by concatenated segment GI is crossed, this is an additional logical relationship crossed. This is because, from the root of the logical database, the logical relationship made possible by concatenated segment type BF must be crossed to allow access to concatenated segment GI.

Rules for Defining Logical Relationships

When the first logical relationship is crossed in a hierarchic path of a logical database, access to all segments in the hierarchic path of the destination parent is made possible as follows:

- Parent segments of the destination parent are included in the logical database as dependents of the destination parent in reverse order, as shown in Figure 104.
- Dependent segments of the destination parent are included in the logical database as dependents of the destination parent without their order changed, as shown in Figure 104.

When an additional logical relationship is crossed in a logical database, access to all segments in the hierarchic path of the destination parent is made possible, just as in the first crossing.

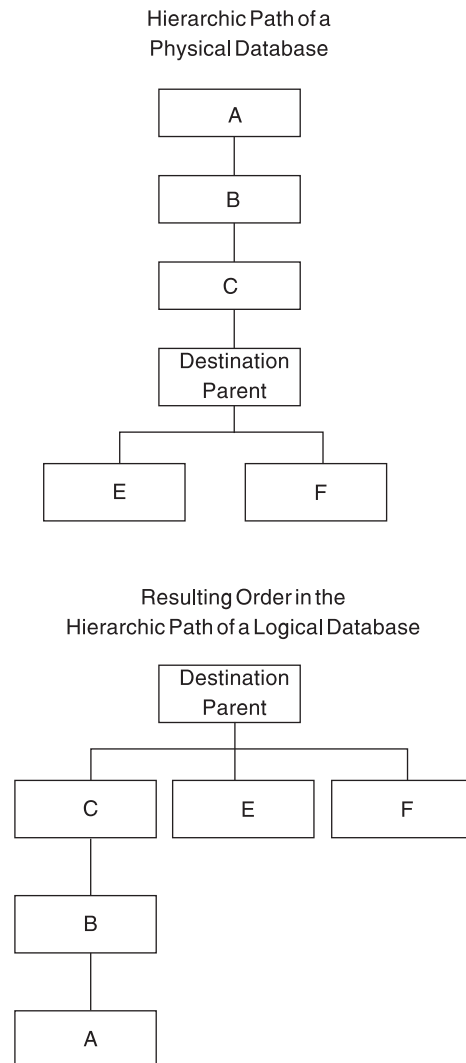


Figure 104. Logical Database Hierarchy Enabled by Crossing the First Logical Relationship

Rules for Defining Logical Databases:

- The root segment in a logical database must be the root segment in a physical database.

- A logical database must use only those segments and physical and logical relationship paths defined in the physical DBD referenced by the logical DBD.
- The path used to connect a parent and child in a logical database must be defined as a physical relationship path or a logical relationship path in the physical DBD referenced by the logical DBD.
- Physical and logical relationship paths can be mixed in a hierarchic segment path in a logical database.
- Additional physical relationship paths, logical relationship paths, or both paths can be included after a logical relationship is crossed in a hierarchic path in a logical database. These paths are included by going in upward directions, downward directions, or both directions, from the destination parent. When proceeding downward along a physical relationship path from the destination parent, direction cannot be changed except by crossing a logical relationship. When proceeding upward along a physical relationship path from the destination parent, direction can be changed.
- Dependents in a logical database must be in the same relative order as they are under their parent in the physical database. If a segment in a logical database is a concatenated segment, the physical children of the logical child and children of the destination parent can be in any order. The relative order of the children or the logical child and the relative order of the children of the destination parent must remain unchanged.
- The same concatenated segment type can be defined multiple times with different combinations of key and data sensitivity. Each must have a distinct name for that view of the concatenated segment. Only one of the views can have dependent segments. Figure 105 shows the four views of the same concatenated segment that can be defined in a logical database. A PCB for the logical database can be sensitive to only one of the views of the concatenated segment type.




Figure 105. Single Concatenated Segment Type Defined Multiple Times with Different Combinations of Key and Data Sensitivity

- LC** Logical child segment type
- DP** Destination parent segment type
- K** KEY sensitivity specified for the segment type
- D** DATA sensitivity specified for the segment type

Choosing Replace, Insert, and Delete Rules for Logical Relationships

You must establish insert, delete, and replace rules when a segment is involved in a logical relationship, because such segments can be updated from two paths: a physical path and a logical path.

Rules for Defining Logical Relationships

Figure 106 and Figure 107 show example insert, delete, and replace rules. Consider the following questions:

1. Should the CUSTOMER segment in Figure 106 be able to be inserted by both its physical and logical paths?
2. Should the BORROW segment be replaceable using only the physical path, or using both the physical and logical paths?
3. If the LOANS segment is deleted using its physical path, should it be erased from the database? Or should it be marked as physically deleted but remain accessible using its logical path?
4. If the logical child segment BORROW or the concatenated segment BORROW/LOANS is deleted from the physical path, should the logical path CUST/CUSTOMER also be automatically deleted? Or should the logical path remain?

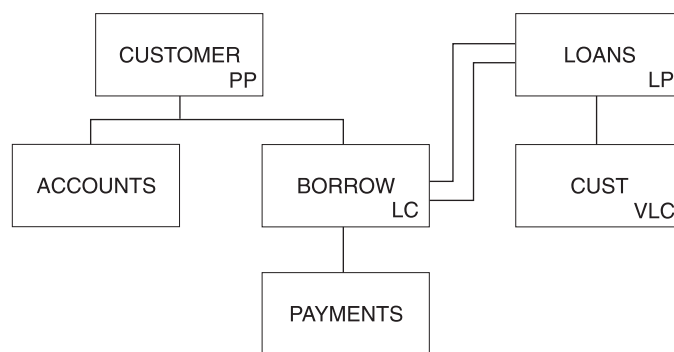


Figure 106. Example of the Replace, Insert, and Delete Rules

| Abbreviation | Explanation |
|--------------|------------------------------------|
| PP | Physical parent segment type |
| LC | Logical child segment type |
| LP | Logical parent segment type |
| VLC | Virtual logical child segment type |

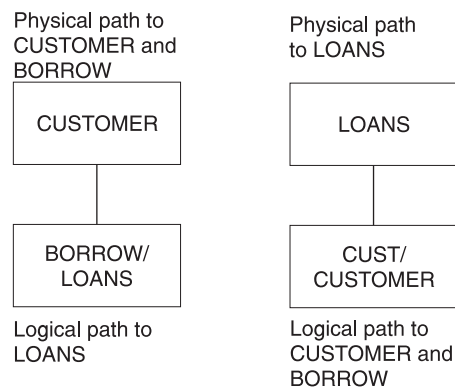


Figure 107. Example of the Replace, Insert, and Delete Rules: Before and After

The answer to these questions depends on the application. The enforcement of the answer depends on your choosing the correct insert, delete, and replace rules for

the logical child, logical parent, and physical parent segments. You must first determine your application processing requirements and then the rules that support those requirements.

For example, the answer to question 1 depends on whether the application requires that a CUSTOMER segment be inserted into the database before accepting the loan. An insert rule of physical (P) on the CUSTOMER segment prohibits insertion of the CUSTOMER segment except by the physical path. An insert rule of virtual (V) allows insertion of the CUSTOMER segment by either the physical or logical path. It probably makes sense for a customer to be checked (past credit, time on current job, and so on.) and the CUSTOMER segment inserted before approving the loan and inserting the BORROW segment. Thus, the insert rule for the CUSTOMER segment should be P to prevent the segment from being inserted logically. (Using the insert rule in this example provides better control of the application.)

Or consider question 3. If it is possible for this loan institution to cancel a type of loan (cancel 10% car loans, for instance, and create 12% car loans) before everyone with a 10% loan has fully paid it, then it is possible for the LOANS segment to be physically deleted and still be accessible from the logical path. This can be done by specifying the delete rule for LOANS as either logical (L) or V, but not as P.

The P delete rule prohibits physically deleting a logical parent segment before all its logical children have been physically deleted. This means the logical path to the logical parent is deleted first.

You need to examine all your application requirements and decide who can insert, delete, and replace segments involved in logical relationships and how those updates should be made (physical path only, or physical and logical path). The insert, delete, and replace rules in the physical DBD and the PROCOPT= parameter in the PCB are the means of control.

Related Reading: These rules are explained in detail in Appendix B, “Insert, Delete, and Replace Rules for Logical Relationships,” on page 471.

Performance Considerations for Logical Relationships

If you are implementing a logical relationship, you make several choices that affect the resources needed to process logically related segments. This topic explains these choices.

Logical Parent Pointers

The logical child segment on DASD has a pointer to its logical parent. You choose how this pointer is physically stored on external storage. Your choices are:

- Direct pointing (specified by coding POINTER=LPARNT in the SEGM statement for the logical child)
- Symbolic pointing (specified by coding the PHYSICAL operand for the PARENT= keyword in the SEGM statement for the logical child)
- Both direct and symbolic pointing

The advantages of direct pointers are:

- Because direct pointers are only 4 bytes long, they are usually shorter than symbolic pointers. Therefore, less DASD space is generally required to store direct pointers.

Performance Considerations for Logical Relationships

- Direct pointers usually give faster access to logical parent segments, except possibly HDAM or PHDAM logical parent segments, which are roots. Symbolic pointers require extra resources to search an index for a HIDAM database. Also, with symbolic pointers, DL/I has to navigate from the root to the logical parent if the logical parent is not a root segment.

The advantages of symbolic pointers are:

- Symbolic pointers are stored as part of the logical child segment on DASD. Having the symbolic key stored on DASD can save the resources required to format a logical child segment in the user's I/O area. Remember, the symbolic key always appears in the I/O area as part of the logical child. When retrieving a logical child, IMS has to construct the symbolic key if it is not stored on DASD.
- Logical parent databases can be reorganized without the logical child database having to be reorganized. This applies to unidirectional and bidirectional physically paired relationships (when symbolic pointing is used).

Symbolic pointing must be used:

- When pointing to a HISAM logical parent database
- If you need to sequence logical child segments (except virtual logical children) on any part of the symbolic key

KEY/DATA Considerations

When you include a concatenated segment as part of a logical DBD, you control how the concatenated segment appears in the user's I/O area. You do this by specifying either KEY or DATA on the SOURCE= keyword of the SEGM statement for the concatenated segment. A concatenated segment consists of a logical child followed by a logical (or destination) parent. You specify KEY or DATA for both parts. For example, you can access a concatenated segment and ask to see the following segment parts in the I/O area:

- The logical child part only
- The logical (or destination) parent part only
- Both parts

By carefully choosing KEY or DATA, you can retrieve a concatenated segment with fewer processing and I/O resources. For example:

- Assume you have the unidirectional logical relationship shown in Figure 108 on page 191.

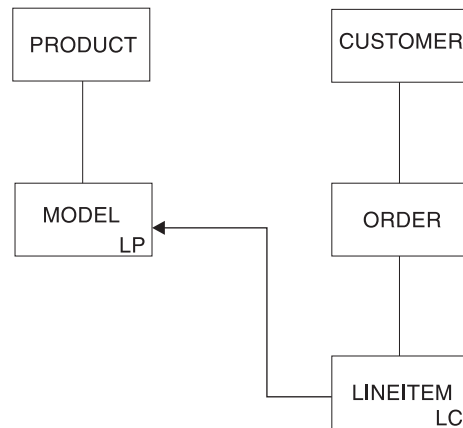


Figure 108. Example of a Unidirectional Logical Relationship

- Assume you have the logical structure shown in Figure 109.




Figure 109. Example of a Logical Structure

- Finally, assume you only need to see the data for the LINEITEM part of the concatenated segment.

You can avoid the extra processing and I/O required to access the MODEL part of the concatenated segment if you:

- Code the SOURCE keyword of the concatenated segment's SEGM statement as:
`SOURCE=(1csegname,DATA,1cdbname),(1psegname,KEY,1pdbname)`
- Store a symbolic logical parent pointer in LINEITEM. If you do not store the symbolic pointer, DL/I must access MODEL and PRODUCT to construct it.

To summarize, do not automatically choose DATA sensitivity for both the logical child and logical parent parts of a concatenated segment. If you do not need to see the logical parent part, code KEY sensitivity for the logical parent and store the symbolic logical parent pointer on DASD.

Sequencing Logical Twin Chains

With virtual pairing, it is possible to sequence the real logical child on physical twin chains and the virtual logical child on logical twin chains. If possible, avoid operations requiring that you sequence logical twins. When a logical twin chain is

Performance Considerations for Logical Relationships

followed, DL/I usually has to access multiple database records. Accessing multiple database records increases the resources required to process the call.

This problem of increased resource requirements to process calls is especially severe when you sequence the logical twin chain on all or part of the symbolic logical parent pointer. Because a virtual logical child is not stored, it is necessary to construct the symbolic logical parent pointer to determine if a virtual logical child satisfies the sequencing operation. DL/I must follow physical parent pointers to construct the symbolic pointers. This process takes place for each virtual logical child in the logical twin chain until the correct position is found for the sequencing operation.

Placement of the Real Logical Child in a Virtually Paired Relationship

In placing the real logical child in a virtually paired relationship, here are some considerations:

- If you need the logical child sequenced in only one of the logically related databases, put the real logical child in that database.
- If you must sequence the logical child in both logically related databases, put the real logical child in the database from which it is most often retrieved.
- Try to place the real logical child so logical twin chains are as short as possible. This placement decreases the number of database records that must be examined to follow a logical twin chain.

Note: You cannot store a real logical child in a HISAM database, because you cannot have logical child pointers (which are direct pointers) in a HISAM database.

Secondary Indexes

The following database types support secondary indexes:

- HISAM
- SHISAM
- HDAM
- PHDAM
- HIDAM
- PHIDAM

Secondary indexes are indexes that allow you to process a segment type in a sequence other than the one defined by the segment's key. A secondary index can also be used to process a segment type based on a qualification in a dependent segment.

Why Secondary Indexes?

When you design your database records, you design them to meet the processing requirements of many applications. You decide what segments will be in a database record and what fields will be in a segment. You decide the order of segments in a database record and fields within a segment. You also decide which field in the root segment will be the key field, and whether the key field will be unique. All these decisions are based on what works best for all your application's processing requirements. However, the choices you make might suit the processing requirements of some applications better than others.

Example: A database record in an educational database is shown in Figure 110.

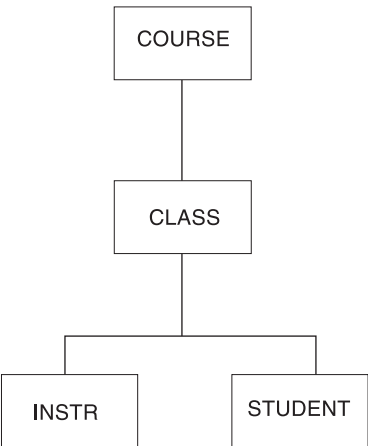


Figure 110. Database Record in Educational Database

Figure 111, shows the root segment, COURSE, and the fields it contains. The course number field is a unique key field.

| | | | | | |
|------------|---------------|-------------|-------------------|-----------|----------------|
| Class Date | Course Number | Course Name | Class Room Number | Room Size | Total Attended |
| | Key Field | | | | |

Figure 111. Example of a Database Record Unique Key Field

You chose COURSE as the root and course number as a unique key field partly because most applications requested information based on course numbers. For these applications, access to the information needed from the database record is fast. For a few of your applications, however, the organization of the database record does not provide such fast access. One application, for example, would be to access the database by student name and then get a list of courses a student is taking. Given the order in which the database record is now organized, access to the courses a student is taking requires a sequential scan of the entire database. Each database record has to be checked for an occurrence of the STUDENT segment. When a database record for the specific student is found, then the COURSE segment has to be referenced to get the name of the course the student is taking. This type of access is relatively slow. In this situation, you can use a secondary index that has a set of pointer segments for each student to all COURSE segments for that student.

Another application would be to access COURSE segments by course name. In this situation, you can use a secondary index that allows access to the database in course name sequence (rather than by course number, which is the key field).

Secondary indexing is a solution to the different processing requirements of various applications. It allows you to have an index based on any field in the database, and not just the key field in the root segment.

Characteristics of Secondary Indexes

Secondary indexes can be used with HISAM, HDAM, PHDAM, HIDAM, and PHIDAM databases. A secondary index is in its own separate database and must use VSAM as its access method. Because a secondary index is in its own database, it can be processed as a separate database.

Secondary indexes are invisible to the application program. When an application program needs to do processing using the secondary index, this fact is communicated to IMS by coding the PROCSEQ= parameter in the PCB. If an application program needs to do processing using the regular processing sequence, PROCSEQ= is simply not coded. If the application program needs to do processing using both the regular processing sequence and the secondary index, the application program's PSB must contain two PCBs, one with PROCSEQ= coded and one without.

When two PCBs are used, it enables an application program to use two paths into the database and two sequence fields. One path and sequence field is provided by the regular processing sequence, and one is provided by the secondary index. The secondary index gives an application program both an alternative way to enter the database and an alternative way to sequentially process database records.

A final characteristic of secondary indexes is that there can be 32 secondary indexes for a segment type and a total of 1000 secondary indexes for a single database.

Segments Used for Secondary Indexes

As shown in Figure 112, to set up a secondary index, three types of segments must be defined to IMS. The three types of segments are pointer, target, and source segments.

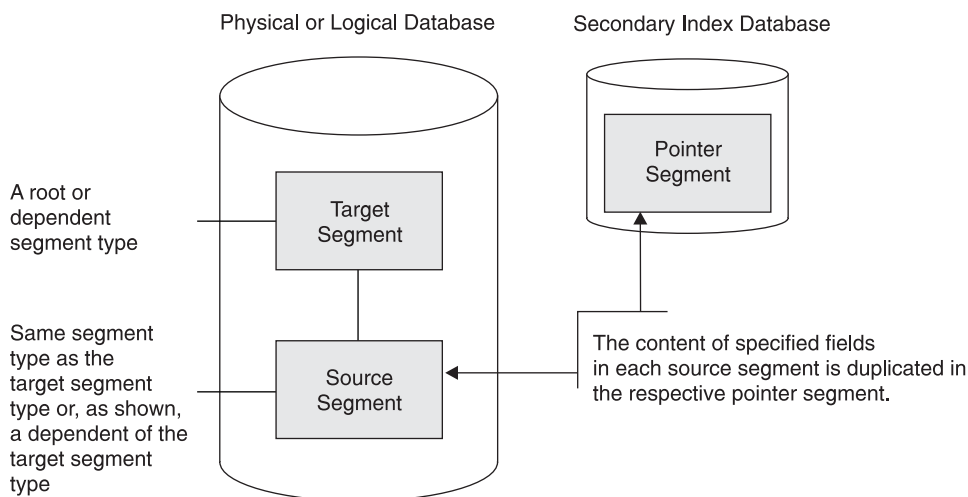


Figure 112. Segments Used for Secondary Indexes

- **Pointer Segment.** The pointer segment is contained in the secondary index database and is the *only* type of segment in the secondary index database. Its format is shown in Figure 113 on page 195.




Figure 113. Format of Pointer Segments Contained in the Secondary Index Database

The first field in the prefix is the delete byte. The second field is the address of the segment the application program retrieves from the regular database. This field is not present if the secondary index uses symbolic pointing. Symbolic pointing is pointing to a segment using its concatenated key. HIDAM and HDAM can use symbolic pointing; however, HISAM must use symbolic pointing. Symbolic pointing is not supported for PHDAM and PHIDAM databases.

For a HALDB PSINDEX database, the segment prefix of pointer segments is slightly different. The “RBA of the segment to be retrieved field” is part of an Extended Pointer Set (EPS), which is longer than 4 bytes. Within the prefix the EPS is followed by the key of the target’s root.

- **Target Segment.** The target segment is in the regular database, and it is the segment the application program needs to retrieve. A target segment is the segment to which the pointer segment points. The target segment can be at any one of the 15 levels in the database, and it is accessed directly using the RBA or symbolic pointer stored in the pointer segment. Physical parents of the target segment are not examined to retrieve the target segment (except in one special case discussed in “Concatenated Key Field” on page 201).
- **Source Segment.** The source segment is also in the regular database. The source segment contains the field (or fields) that the pointer segment has as its key field. Data is copied from the source segment and put in the pointer segment’s key field. The source and the target segment can be the same segment, or the source segment can be a dependent of the target segment. The optional fields are also copied from the source segment with one exception, which is discussed later in this topic.

Using the education database in Figure 114 on page 196, you can see how three segments work together. In this example, the education database is a HIDAM database that uses RBAs rather than symbolic pointers. Suppose an application program needs to access the education database by student name and then list all courses the student is taking:

- The segment the application is trying to retrieve is the COURSE segment, because the segment contains the names of courses (COURSENM field). Therefore, COURSE is the target segment, and needs retrieval.
- In this example, the application program is going to use the student’s name in its DL/I call to retrieve the COURSE segment. The DL/I call is qualified using student name as its qualifier. The source segment contains the fields used to sequence the pointer segments in the secondary index. In this example, the pointer segments must be sequenced by student name. The STUDENT segment becomes the source segment. It is the fields in this segment that are copied into the data portion of the pointer segment as the key field.
- The call from the application program invokes a search for a pointer segment with a key field that matches the student name. Once the correct pointer segment in the index is found, it contains the address of the COURSE segment the application program is trying to retrieve.

Secondary Indexes




Figure 114. Education Database Record

Figure 115 shows how the pointer, target, and source segments work together. Figure 116 is the call the application program issues. XNAME is the from the NAME parameter in the XFLD statement.




Figure 115. How a Segment Is Accessed Using a Secondary Index

```
GU COURSE (XNAME = BAKER ... )
```

Figure 116. Call an Application Issues When a Secondary Index Is Used

COURSE is the target segment that the application program is trying to retrieve.

STUDENT is the source segment containing the one or more fields that the application program uses as a qualifier in its call and that the data portion of a pointer segment contains as a key.

The BAKER segment in the secondary index is the pointer segment, whose prefix contains the address of the segment to be retrieved and whose data fields contain the key the application program uses as a qualifier in its call.

How the Hierarchy Is Restructured

When the PROCSEQ= parameter in the PCB is coded (specifying that the application program needs to do processing using the secondary index), the way in which the application program perceives the database record changes.

If the target segment is the root segment in the database record, the structure the application program perceives does not differ from the one it can access using the regular processing sequence. However, if the target segment is not the root segment, the hierarchy in the database record is conceptually restructured. Figure 117 and Figure 118 on page 198 illustrate this concept.

The target segment (as shown in the figure) is segment G. Target segment G becomes the root segment in the restructured hierarchy. All dependents of the target segment (segments H, J, and I) remain dependents of the target segment. However, all segments on which the target is dependent (segments D and A) and their subordinates become dependents of the target and are put in the left most positions of the restructured hierarchy. Their position in the restructured hierarchy is the order of immediate dependency. D becomes an immediate dependent of G, and A becomes an immediate dependent of D.




Figure 117. Physical Database Structure with Target Segment G




Figure 118. Secondary Index Structure Indexed in Secondary Index on Segment G

Secondary Data Structure

This new structure is called a *secondary data structure*. A processing restriction exists when using a secondary data structure, and the target segment and the segments on which it was dependent (its physical parents, segments D and A) cannot be inserted or deleted.

Secondary Processing Sequence

The restructuring of the hierarchy in the database record changes the way in which the application program accesses segments. The new sequence in which segments are accessed is called the *secondary processing sequence*. Figure 118 shows how the application program perceives the database record.

If the same segment is referenced more than once (as shown in Figure 118), you must use the DBDGEN utility to generate a logical DBD that assigns alternate names to the additional segment references. If you do not generate the logical DBD, the PSBGEN utility will issue the message “SEG150” for the duplicate SENSEG names.

How a Secondary Index Is Stored

Secondary index databases contain root segments only. They are stored in a single VSAM KSDS if the key in the pointer segment is unique. If keys are not unique, an additional data set must be used (an ESDS) to store segments containing duplicate keys. (KSDS data sets do not allow duplicate keys.) Duplicate keys exist when, for example, a secondary index is used to retrieve courses based on student name. As shown in Figure 119 on page 199, several source segments could exist for each student.

| Prefix | | Data | |
|--------|-----|------|-------|
| DB | RBA | MATH | ADAMS |

| | | | |
|----|-----|--------|-------|
| DB | RBA | FRENCH | ADAMS |
|----|-----|--------|-------|

| | | | |
|----|-----|------|-------|
| DB | RBA | HIST | ADAMS |
|----|-----|------|-------|

Figure 119. Examples of Source Segments for Each Student

Each pointer segment in a secondary index is stored in one logical record. A logical record containing a pointer segment is shown in Figure 120.




Figure 120. Example of a Logical Record Containing a Pointer Segment

A HALDB secondary index entry is shown in Figure 121.

| Delete Byte | EPS-Extended pointer set | Key for the root of the target | Data |
|-------------|--------------------------|--------------------------------|-----------------|
| 1 | 26 | 1 to 255 | Variable length |

Figure 121. Secondary Index Entry for HALDB

The format of the logical record is the same in both a KSDS and ESDS data set. The pointer field at the beginning of the logical record exists only when the key in the data portion of the segment is not unique. If keys are not unique, some pointer segments will contain duplicate keys. These pointer segments must be chained together, and this is done using the pointer field at the beginning of the logical record.

Pointer segments containing duplicate keys are stored in the ESDS in LIFO (last in, first out) sequence. When the first duplicate key segment is inserted, it is written to the ESDS, and the KSDS logical record containing the segment it is a duplicate of points to it. When the second duplicate is inserted, it is inserted into the ESDS in the next available location. The KSDS logical record is updated to point to the second duplicate. The effect of inserting duplicate pointer segments into the ESDS in LIFO sequence is that the original pointer segment (the one in the KSDS) is retrieved last. This retrieval sequence should not be a problem, because duplicates, by definition, have no special sequence.

Format and Use of Fields in a Pointer Segment

This topic contains diagnosis, modification, or tuning information.

Figure 122 on page 200 shows the fields in a pointer segment. Like all segments, the pointer segment has a prefix and data portion. The prefix portion has a delete byte, and when direct rather than symbolic pointing is used, it has the address of the target segment (4 bytes). The data portion has a series of fields, and some of them are optional. All fields in the data portion of a pointer segment contain data

Secondary Indexes

taken from the source segment (with the exception of user data). These fields are the constant field (optional), the search field, the subsequence field (optional), the duplicate data field (optional), the concatenated key field (optional except for HISAM), and then the data (optional).

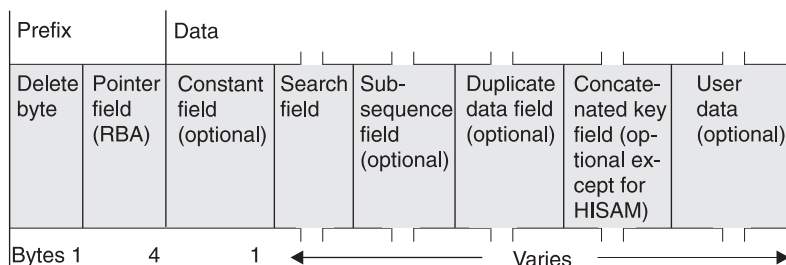


Figure 122. Examples of Several Source Segments for Each Student

Delete Byte

The delete byte is used by IMS to determine whether a segment has been deleted from the database.

Pointer Field

This field, when present, contains the RBA of the target segment. The pointer field exists when direct pointing is specified for an index pointing to an HD database. Direct pointing is simply pointing to a segment using its actual address. The other type of pointing that can be specified is symbolic pointing. Symbolic pointing, which is explained under “Concatenated Key Field,” can be used to point to HD databases and must be used to point to HISAM databases. If symbolic pointing is used, this field does not exist.

Constant Field

This field, when present, contains a 1-byte constant. The constant is used when more than one index is put in an index database (This topic is discussed under “Sharing Secondary Index Databases” on page 207). The constant identifies all pointer segments for a specific index in the shared index database. The value in the constant field becomes part of the key.

Search Field

The data in the search field is the key of the pointer segment. All data in the search field comes from data in the source segment. As many as five fields from the source segment can be put in the search field. These fields do not need to be contiguous fields in the source segment. When the fields are stored in the pointer segment, they can be stored in any order. When stored, the fields are concatenated. The data in the search field (the key) can be unique or non-unique.

IMS automatically maintains the search field in the pointer segment whenever a source segment is modified.

Subsequence Field

The subsequence field, like the search field, contains from one to five fields of data from the source segment. Subsequence fields are optional, and can be used if you have non-unique keys. The subsequence field can make non-unique keys unique. Making non-unique keys unique is desirable because of the many disadvantages of non-unique keys. For example, non-unique keys require you to use an additional data set, an ESDS, to store all index segments with duplicate keys. An ESDS requires additional space. More important, the search for specific occurrences of duplicates requires additional I/O operations that can decrease performance.

When a subsequence field is used, the subsequence data is concatenated with the data in the search field. These concatenated fields become the key of the pointer segment. If properly chosen, the concatenated fields form a unique key. (It is not always possible to form a unique key using source data in the subsequence field. Therefore, you can use system related fields, explained later in the chapter, to form unique keys.)

One important thing to note about using subsequence fields is that if you use them, the way in which an SSA is coded does not need to change. The SSA can still specify what is in the search field, but it cannot specify what is in the search plus the subsequence field. Subsequence fields are not seen by the application program unless it is processing the secondary index as a separate database.

Up to five fields from the source segment can be put in the subsequence field. These fields do not need to be contiguous fields in the source segment. When the fields are stored in the pointer segment, they can be stored in any order. When stored, they are concatenated.

IMS automatically maintains the subsequence field in the pointer segment whenever a source segment is modified.

Duplicate Data Field

The duplicate data field, like the search field, contains from one to five fields of data from the source segment. Duplicate data fields are optional. Use duplicate data fields when you have applications that process the secondary index as a separate database (This topic is discussed under “Processing a Secondary Index as a Separate Database” on page 206). Like the subsequence field, the duplicate data field is not seen by an application program unless it is processing the secondary index as a separate database.

As many as five fields from the source segment can be put in the duplicate data field. These fields do not need to be contiguous fields in the source segment. When the fields are stored in the pointer segment, they can be stored in any order. When stored, they are concatenated.

IMS automatically maintains the duplicate data field in the pointer segment whenever a source segment is modified.

Concatenated Key Field

This field, when present, contains the concatenated key of the target segment. This field exists when the pointer segment points to the target segment symbolically, rather than directly. Direct pointing is simply pointing to a segment using its actual address. Symbolic pointing is pointing to a segment by a means other than its actual address. In a secondary index, the concatenated key of the target segment is used as a symbolic pointer.

Segments in an HDAM or a HIDAM database being accessed using a secondary index can be accessed using a symbolic pointer. Segments in a HISAM database must be accessed using a symbolic pointer because segments in a HISAM database can “move around,” and the maintenance of direct-address pointers could be a large task. One of the implications of using symbolic pointers is that the physical parents of the target segment must be accessed to get to the target segment. Because of this extra access, retrieval of target segments using symbolic pointing is not as fast as retrieval using direct pointing. Also, symbolic pointers generally require more space in the pointer segment. When symbolic pointers are

Secondary Indexes

used, the pointer field (4 bytes long) in the prefix is not present, but the fully concatenated key of the target segment is generally more than 4 bytes long.

IMS automatically generates the concatenated key field when symbolic pointing is specified.

One situation exists in which symbolic pointing is specified and IMS does not automatically generate the concatenated key field. This situation is caused by specifying the system-related field /CK as a subsequence or duplicate data field in such a way that the concatenated key is fully contained. In this situation, the symbolic pointer portion of either the subsequence field or the duplicate data field is used.

User Data in Pointer Segments

You can include any user data in the data portion of a pointer segment by specifying a segment length long enough to hold it. You need user data when applications process the secondary index as a separate database (This topic is discussed under “Processing a Secondary Index as a Separate Database” on page 206). Like data in the subsequence and duplicate data fields, user data is never seen by an application program unless it is processing the secondary index as a separate database.

You must initially load user data. You must also maintain it. During reorganization of a database that uses secondary indexes, the secondary index database is rebuilt by IMS. During this process, all user data in the pointer segment is lost.

Making Keys Unique Using System Related Fields

You have already seen why it is desirable to have unique keys in the secondary index. You have also seen one way to force unique keys using the subsequence field in the pointer segment. If use of the subsequence field to contain additional information from the source segment does not work for you, there are two other ways to force unique keys. Both are done using an operand in the FIELD statement of the source segment in the DBD. The FIELD statement defines fields within a segment type.

Using the /SX Operand

For HD databases, you can code a FIELD statement with a NAME field that starts with /SX. The /SX can be followed by any additional characters (up to five) that you need. When you use this operand, the system generates (during segment insertion) the RBA, or an 8-byte ILK for PHDAM or PHIDAM, of the source segment. The system also puts the RBA or ILK in the subsequent field in the pointer segment, thus ensuring that the key is unique. The FIELD statement in which /SX is coded is the FIELD statement defining fields in the source segment. The /SX value is not, however, put in the source segment. It is put in the pointer segment.

When you use the /SX operand, the XDFLD statement in the DBD must also specify /SX (plus any of the additional characters added to the /SX operand). The XDFLD statement, among other things, identifies fields from the source segment that are to be put in the pointer segment. The /SX operand is specified in the SUBSEQ= operand in the XDFLD statement.

Using the /CK Operand

The other way to force unique keys is to code a FIELD statement with a NAME parameter that starts with /CK. When used as a subsequence field, /CK ensures unique keys for pointer segments. You can use this operand for HISAM, HDAM, PHDAM, HIDAM, or PHIDAM databases. The /CK can be followed by up to five

additional characters. The /CK operand works like the /SX operand except that the concatenated key, rather than the RBA, of the source segment is used. Another difference is that the concatenated key is put in the subsequence or duplicate data field in the pointer segment. Where the concatenated key is put depends on where you specify the /CK.

When using /CK, you can use a portion of the concatenated key of the source segment (if some portion will make the key unique) or all of the concatenated key. You use the BYTES= and START= operands in the FIELD statement to specify what you need.

For example, suppose you are using the database record shown in Figure 123.

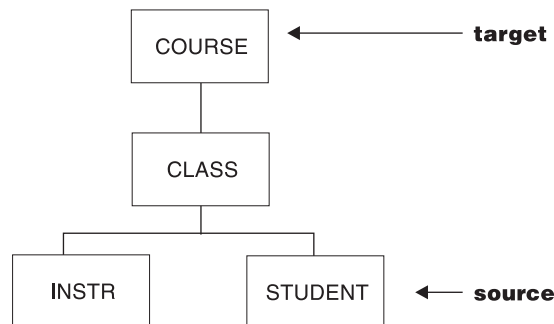


Figure 123. Database Record Showing the Source and Target for Secondary Indexes

The concatenated key of the STUDENT segment is shown in Figure 124.

| COURSECD | CLASSNO | SEQ |
|----------|---------|-----|
| Bytes 15 | 3 | 3 |

Figure 124. Concatenated Key of the STUDENT Segment

If you specify on the FIELD statement whose name begins with /CK BYTES=21, START=1, the entire concatenated key of the source segment will be put in the pointer segment. If you specify BYTES=6, START=16, only the last six bytes of the concatenated key (CLASSNO and SEQ) will be put in the pointer segment. The BYTES= operand tells the system how many bytes are to be taken from the concatenated key of the source segment in the PCB key feedback area. The START= operand tells the system the beginning position (relative to the beginning of the concatenated key) of the information that needs to be taken. As with the /SX operand, the XDFLD statement in the DBD must also specify /CK.

To summarize: /SX and /CK fields can be included on the SUBSEQ= parameter of the XDFLD statement to make key fields unique. Making key fields unique avoids the overhead of using an ESDS to hold duplicate keys. The /CK field can also be specified on the DDATA= parameter of the XDFLD statement but the field will not become part of the key field.

When making keys unique, unique sequence fields must be defined in the target segment type, if symbolic pointing is used. Also, unique sequence fields must be defined in all segment types on which the target segment type is dependent (in the physical rather than restructured hierarchy in the database).

Suppressing Index Entries: Sparse Indexing

When a source segment is loaded, inserted, or replaced in the database, DL/I automatically creates or maintains the pointer segment in the index. This happens automatically unless you have specified that you do not need certain pointer segments built.

For example: suppose you have a secondary index for the education database at which you have been previously looking. STUDENT is the source segment, and COURSE is the target segment. You might need to create pointer segments for students only if they are associated with a certain customer number. This could be done using sparse indexing, a performance enhancement of secondary indexing.

Advantages of Sparse Indexing

Sparse indexing allows you to specify the conditions under which a pointer segment is suppressed, not generated, and put in the index database. Sparse indexing has two advantages. The primary one is that it reduces the size of the index, saving space and decreasing maintenance of the index. By decreasing the size of the index, performance is improved. The second advantage is that you do not need to generate unnecessary index entries.

How to Specify a Sparse Index

Sparse indexing can be specified in two ways:

- You can code a value in the NULLVAL= operand on the XDFLD statement in the DBD that equals the condition under which you do not need a pointer segment put in the index. You can put BLANK, ZERO, or any 1-byte value (for example, X'10', C'Z', 5, or B'00101101') in the NULLVAL= operand.
 - BLANK is the same as C ' ' or X'40'
 - ZERO is the same as X'00' but *not* C'0'

When using the NULLVAL= operand, a pointer segment is suppressed if every byte of the source field has the value you used in the operand.

- If the values you are allowed to code in the NULLVAL= operand do not work for you, you can create an index maintenance exit routine that determines the condition under which you do not need a pointer segment put in the index. If you create your own index maintenance exit routine, you code its name in the EXTRTN= operand on the XDFLD statement in the DBD. You can only have one index maintenance exit routine for each secondary index. This exit routine, however, can be a general purpose one that is used by more than one secondary index.

The exit routine must be consistent in determining whether a particular pointer segment needs to be put in the index. The exit routine cannot examine the same pointer segment at two different times but only mark it for suppression once. Also, user data cannot be used by your exit routine to determine whether a pointer segment is to be put in the index. When a pointer segment needs to be inserted into the index, your exit routine only sees the actual pointer segment just before insertion. When a pointer segment is being replaced or deleted, only a prototype of the pointer segment is seen by your exit routine. The prototype contains the contents of the constant, search, subsequence, and duplicate data fields, plus the symbolic pointer if there is one.

The information needed to code a secondary index maintenance exit routine is in *IMS Version 9: Customization Guide*.

How the Secondary Index Is Maintained

When a source segment is inserted, deleted, or replaced in the database, IMS keeps the index current regardless whether the application program performing the update uses the secondary index.

The way in which IMS maintains the index depends on the operation being performed. Regardless of the operation, IMS always begins index maintenance by building a pointer segment from information in the source segment that is being inserted, deleted, or replaced. (This pointer segment is built but not yet put in the secondary index database.)

Inserting a Source Segment

When a source segment is inserted, DL/I determines whether the pointer segment needs to be suppressed. If the pointer segment needs to be suppressed, it is not put in the secondary index. If the pointer segment does not need to be suppressed, it is put in the secondary index.

Deleting a Source Segment

When a source segment is deleted, IMS determines whether the pointer segment is one that was suppressed. If so, IMS does not do any index maintenance. If the segment is one that was suppressed, there should not be a corresponding pointer segment in the index to delete. If the pointer segment is not one that was suppressed, IMS finds the matching pointer segment in the index and deletes it. Unless the segment contains a pointer to the ESDS data set, which can occur with a non-unique secondary index, the logical record containing the deleted pointer segment in a KSDS data set is erased.

Replacing a Source Segment

When a source segment is replaced, the pointer segment in the index might or might not be affected. The pointer segment in the index might need to be replaced, or it might need to be deleted. After replacement or deletion, a new pointer segment is inserted. On the other hand, the pointer segment might need no changes. IMS determines what needs to be done by comparing the pointer segment it built (the new one) with the matching pointer segment in the secondary index (the old one).

- If both the new and the old pointer segments need to be suppressed, IMS does not do anything (no pointer segment exists in the index).
- If the new pointer segment needs to be suppressed but the old one does not, then the old pointer segment is deleted from the index.
- If the new pointer segment does not need to be suppressed but the old pointer segment is suppressed, then the new pointer segment is inserted into the secondary index.
- If neither the new or the old segment needs to be suppressed and:
 - If there is no change to the old pointer segment, IMS does not do anything.
 - If the non-key data portion in the new pointer segment is different from the old one, the old pointer segment is replaced. User data in the index pointer segment is preserved when the pointer segment is replaced.
 - If the key portion in the new pointer segment is different from the old one, the old pointer segment is deleted and the new pointer segment is inserted. User data is *not* preserved when the index pointer segment is deleted and a new one inserted.

If you reorganize your secondary index and it contains non-unique keys, the resulting pointer segment order can be unpredictable.

Processing a Secondary Index as a Separate Database

Because they are actual databases, secondary indexes can be processed independently. A number of reasons exist why an application program might process a secondary index as an independent database. For example, an application program can use the secondary index to retrieve a small piece of data from the database. If you put this piece of data in the pointer segment, the application program can retrieve it without an I/O operation to the regular database. You could put the piece of data in the duplicate data field in the pointer segment if the data was in the source segment. Otherwise, you must carry the data as user data in the pointer segment. (If you carry the data as user data, it is lost when the primary database is reorganized and the secondary index is recreated.)

Another reason for processing a secondary index as a separate database is to maintain it. You could, for example, scan the subsequence or duplicate data fields to do logical comparisons or data reduction between two or more indexes. Or you can add to or change the user data portion of the pointer segment. The only way an application program can see user data or the contents of the duplicate data field is by processing the secondary index as a separate database.

In processing a secondary index as a separate database, several processing restrictions designed primarily to protect the secondary index database exist. The restrictions are as follows:

- Segments cannot be inserted.
- Segments can be deleted. Note, however, that deleted segments can make your secondary index invalid for use as an index.
- The key field in the pointer segment (which consists of the search field, and if they exist, the constant and subsequence fields) cannot be replaced.

In addition to the restrictions imposed by the system to protect the secondary index database, you can further protect it using the PROT operand in the DBD statement. When PROT is specified, an application program can only replace user data in a pointer segment. However, pointer segments can still be deleted when PROT is specified. When a pointer segment is deleted, the source segment that caused the pointer segment to be created is not deleted. Note the implication of this: IMS might try to do maintenance on a pointer segment that has been deleted. When it finds no pointer segment for an existing source segment, it will return an NE status code. When NOPROT is specified, an application program can replace all fields in a pointer segment except the constant, search, and subsequence fields. PROT is the default for this parameter.

For an application program to process a secondary index as a separate database, you merely code a PCB for the application program. This PCB must reference the DBD for the secondary index. When an application program uses qualified SSAs to process a secondary index database, the SSAs must use the complete key of the pointer segment as the qualifier. The complete key consists of the search field and the subsequence and constant fields (if these last two fields exist). The PCB key feedback area in the application program will contain the entire key field.

If you are using a shared secondary index, calls issued by an application program (for example, a series of GN calls) will not violate the boundaries of the secondary index they are against. Each secondary index in a shared database has a unique DBD name and root segment name.

Sharing Secondary Index Databases

As many as 16 secondary indexes can be put in a single index database. When more than one secondary index is in the same database, the database is called a shared index database.

HALDB does not support shared secondary indexes.

Although using a shared index database can save some main storage, the disadvantages of using a shared index database generally outweigh the small amount of space that is saved by its use. For example, performance can decrease when more than one application program simultaneously uses the shared index database. (Search time is increased because the arm must move back and forth between more than one secondary index.) In addition, maintenance, recovery, and reorganization of the shared index database can decrease performance because all secondary indexes are, to some extent, affected if one is. For example, when a database that is accessed using a secondary index is reorganized, IMS automatically builds a new secondary index. This means all other indexes in the shared database must be copied to the new shared index.

If you are using a shared index database, you need to know the following information:

- A shared index database is created, accessed, and maintained just like an index database with a single secondary index.
- The various secondary indexes in the shared index database do not need to index the same database.
- One shared index database could contain all secondary indexes for your installation (if the number of secondary indexes does not exceed 16).

In a shared index database:

- All index segments must be the same length.
- All keys must be the same length.
- The offset from the beginning of all segments to the search field must be the same. This means all keys must be either unique or non-unique. With non-unique keys, a pointer field exists in the target segment. With unique keys, it does not. So the offset to the key field, if unique and non-unique keys were mixed, would differ by 4 bytes.

If the search fields in your secondary indexes are not the same length, you might be able to force key fields of equal length by using the subsequence field. You can put the number of bytes you need to make each key field an equal length in the subsequence field.

- Each shared secondary index requires a constant specified for it, a constant that uniquely identifies it from other indexes in the secondary index database. IMS puts this identifier in the constant field of each pointer segment in the secondary index database. For shared indexes, the key is the constant, search, and (if used) the subsequence field.

Using the INDICES= Parameter

In the PCB on a SENSEG statement, you can specify an INDICES= parameter. This parameter is used to specify a secondary index that contains search fields used to qualify SSAs for an indexed segment type. Figure 125, Figure 126 on page 208, and Figure 127 on page 208 illustrate the use of the INDICES=parameter.

Secondary Indexes

The use of the `INDICES=` parameter does not alter the processing sequence selected for the PCB by the presence or absence of the `PROCSEQ=` parameter.

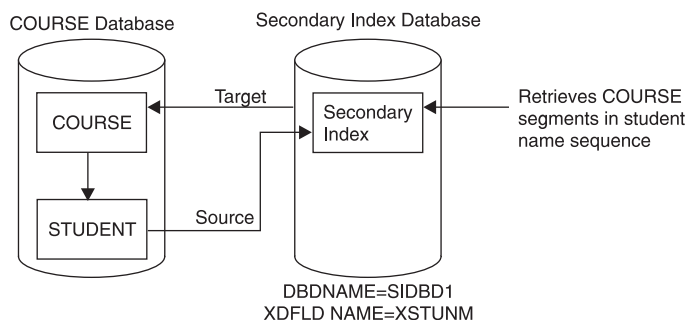


Figure 125. Databases for First Example of the `INDICES=` Parameter

```
PCB
SENSEG NAME=COURSE, INDICES=SIDBD1
SENSEG NAME=STUDENT
```

Figure 126. PCB for the First Example of the `INDICES=` Parameter

```
GU COURSE COURSENM=12345&.XSTUNM=JONES
```

Figure 127. Application Program Call Issued for the First Example of the `INDICES=` Parameter

When the call shown in Figure 127 is used, IMS gets the `COURSE` segment with a number 12345. Then IMS gets a secondary index entry, one in which `XSTUNM` is equal to `JONES`. IMS checks to see if the pointer in the secondary index points to the `COURSE` segment with course number 12345. If it does, IMS returns the `COURSE` segment to the application program's I/O area. If the secondary index pointer does not point to the `COURSE` segment with course number equal to 12345, IMS checks for other secondary index entries with `XSTUNM` equal to `JONES` and repeats the compare.

If all secondary index entries with `XSTUNM` equal to `JONES` result in invalid compares, no segment is returned to the application program. By doing this, IMS need not search the `STUDENT` segments for a student with `NAME` equal to `JONES`. This technique involving use of the `INDICES=` parameter is useful when source and target segments are different.

Compare Process and Performance

Excluding `COURSENM=12345` (in Figure 127) from a `GU` call, impacts performance. IMS retrieves the first `COURSE` segment in the `COURSE` database, and then a secondary index entry in which `XSTUNM` is equal to `JONES`. IMS checks to see if the pointer in the secondary index points to the `COURSE` segment just retrieved. If it does, IMS returns the `COURSE` segment to the application program's I/O area. If the secondary index pointer does not point to this `COURSE` segment, IMS checks for other secondary index entries with `XSTUNM` equal to `JONES` and repeats the compare. If all secondary index entries with `XSTUNM` equal to `JONES` result in invalid compares, IMS retrieves the next `COURSE` segment and the secondary index entries as before, then repeats the compare. If all the `COURSE` segments result in invalid compares, no segment is returned to the application program.

The INDICES= parameter can also be used to reference more than one secondary index in the source call. Figure 130 on page 210 shows the use of INDICES=parameter.

In the Figure 128, IMS uses the SIDBD2 secondary index to get the COURSE segment for MATH. IMS then gets a COURSE segment using the SIDBD1. IMS can then compare to see if the two course segments are the same. If they are, IMS returns the COURSE segment to the application program's I/O area. If the compare is not equal, IMS looks for other SIDBD1 pointers to COURSE segments and repeats the compare operations. If there are still no equal compares, IMS checks for other SIDBD2 pointers to COURSE segments and looks for equal compares to SIDBD1 pointers. If all possible compares result in unequal compares, no segment is returned to the application program.

Note: This compare process can severely degrade performance.

Using Secondary Indexes with Logical Relationships

When creating or using a secondary index for a database that has logical relationships, the following restrictions exist:

- A logical child segment or a dependent of a logical child cannot be a target segment.
- A logical child cannot be used as a source segment; however, a dependent of a logical child can.
- A concatenated segment or a dependent of a concatenated segment in a logical database cannot be a target segment.
- When using logical relationships, no qualification on indexed fields is allowed in the SSA for a concatenated segment. However, an SSA for any dependent of a concatenated segment can be qualified on an indexed field.

Figure 128 shows the databases for the second example of the INDICES parameter. Following the databases is the example PCB in Figure 129 and the application programming call in Figure 130 on page 210.

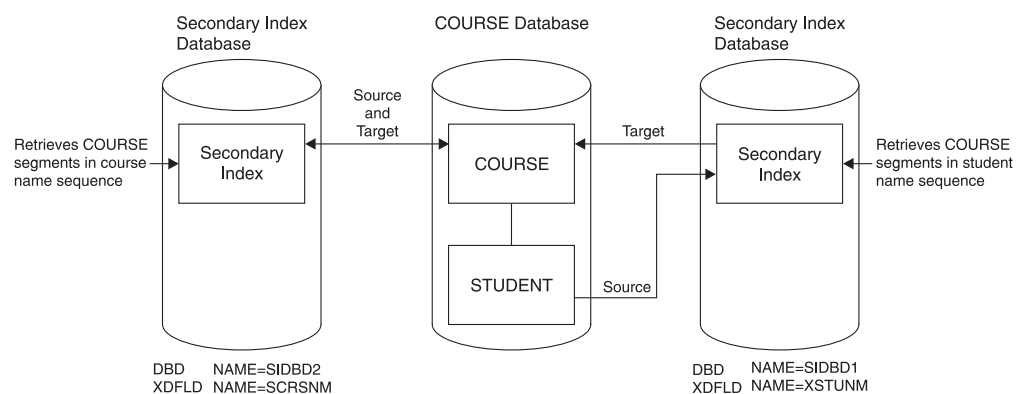


Figure 128. Databases for Second Example of the INDICES= Parameter

```
PCB      PROCSEQ=SIDBD2
SENSEG NAME=COURSE, INDICES=SIDBD1
SENSEG NAME=STUDENT
```

Figure 129. PCB for the Second Example of the INDICES= Parameter

GU COURSE SCRSNM=MATH&XSTUNM=JONES

Figure 130. Application Program Call Issued for the Second Example of the INDICES= Parameter

Using Secondary Indexes with Variable-Length Segments

If a variable-length segment is a source segment, when an occurrence of it is inserted that does not have fields specified for use in the search, subsequence, or duplicate data fields of the pointer segment, the following occurs:

- If the missing source segment data is used in the search field of the pointer segment, no pointer segment is put in the index.
- If the missing source segment data is used in the subsequence or duplicate data fields of the pointer segment, the pointer segment is put in the index. However, the subsequence or duplicate data field will contain one of the three following representations of zero:

P = X'0F'

X = X'00'

C = C'0'

Which of these is used is determined by what is specified on the FIELD statements in the DBD that defined the source segment field.

Considerations When Using Secondary Indexing

- When a source segment is inserted into or deleted from a database, an index pointer segment is inserted into or deleted from the secondary index. This maintenance always occurs regardless of whether the application program doing the updating is using the secondary index.
- When an index pointer segment is deleted by a REPL or DLET call, position is lost for all calls within the database record for which a PCB position was established using the deleted index pointer segment.
- When replacing data in a source segment, if the data is used in the search, subsequence, or duplicate data fields of a secondary index, the index is updated to reflect the change as follows:
 - If data used in the duplicate data field of the pointer segment is replaced in the source segment, the pointer segment is updated with the new data.
 - If data used in the search or subsequence fields of the pointer segment is replaced in the source segment, the pointer segment is updated with the new data. In addition, the position of the pointer segment in the index is changed, because a change to the search or subsequence field of a pointer segment changes the key of the pointer segment. The index is updated by deleting the pointer segment from the position that was determined by the old key. The pointer segment is then inserted in the position determined by the new key.
- The use of secondary indexes increases storage requirements for all calls made within a specific PCB when the processing option allows the source segment to be updated. Additional storage requirements for each secondary index database range from 6K to 10K bytes. Part of this additional storage is fixed in real storage by VSAM.
- You should always compare the use of secondary indexing with other ways of achieving the same result. For example, to produce a report from an HDAM or PHDAM database in root key sequence, you can use a secondary index. However, in many cases, access to each root sequentially is a random operation. It would be very time-consuming to fully scan a large database when access to

each root is random. It might be more efficient to scan the database in physical sequence (using GN calls and no secondary index) and then sort the results by root key to produce a final report in root key sequence.

- When calls for a target segment are qualified on the search field of a secondary index, and the indexed database is not being processed using the secondary index, additional I/O operations are required. Additional I/O operations are required because the index must be accessed each time an occurrence of the target segment is inspected. Because the data in the search field of a secondary index is a duplication of data in a source segment, you should decide whether an inspection of source segments might yield the same result faster.
- When using a secondary data structure, the target segment and the segments on which it was dependent (its physical parents) cannot be inserted or deleted.

How to Specify Use of Secondary Indexing in the DBD

Figure 131 on page 212 shows the EDUC database and its secondary index. Figure 132 on page 213 and Figure 133 on page 213 show the two DBDs required for the databases. The secondary index in this example is used to retrieve COURSE segments based on student names. The example uses direct, rather than symbolic, pointers. The pointer segment in the secondary index contains a student name in the search field and a system related field in the subsequence field. Both of these fields are defined in the STUDENT segment. The STUDENT segment is the source segment. The COURSE segment is the target segment.

The DBDs in Figure 132 on page 213 and Figure 133 on page 213 highlight the statements and parameters coded when a secondary index is used. (Wherever statements or parameters are omitted the parameter in the DBD is coded the same regardless of whether secondary indexing is used.) “DBD for the EDUC Database” and “DBD for the SINDX Database” on page 213 provide a summary of how the statements and parameters in the DBDs in Figure 132 on page 213 and Figure 133 on page 213 are used.

DBD for the EDUC Database

An LCHILD and XDFLD statement are used to define the secondary index. These statements are coded after the SEGM statement for the target segment.

- LCHILD statement. The LCHILD statement specifies the name of the secondary index SEGM statement and the name of the secondary index database in the NAME= parameter. The PTR= parameter is always PTR=INDX when a secondary index is used.
- XDFLD statement. The XDFLD statement defines the contents of the pointer segment and the options used in the secondary index. It must appear in the DBD input deck after the LCHILD statement that references the pointer segment. The meaning of the parameters in the XDFLD statement are as follows:

NAME=

This parameter specifies the name that can be used in the SSA to qualify a DL/I call on the secondary processing sequence.

SEGMENT=

This parameter identifies the source segment, which in this example is STUDENT. If this operand is omitted, the target segment is assumed to be the same segment as the source segment. The remaining parameters in the XDFLD statement describe information related to the source segment.

CONSTANT=

This parameter (not used in the example) specifies the unique constant required when a secondary index is part of a shared database.

Secondary Indexes

SRCH=

This parameter specifies the one to five fields from the source segment that are to be copied into the pointer segment's search field. In this case, only one field is being copied, the STUDNM field, which contains student names.

SUBSEQ=

This parameter specifies the one to five fields from the source segment that are to be copied into the pointer segment's subsequence field. These extra fields can be used to make the key in the index unique. In this case, one field is being copied, the /SX1 field, which contains a system-related field. This parameter is optional.

DDATA=

This parameter (not used in the example) specifies the one to five fields from the source segment that are to be copied into the pointer segment's duplicate data field. These fields can only be accessed when the secondary index is processed as a separate database. This parameter is optional.

NULLVAL=

This parameter (not used in the example) contains a 1-byte value used to suppress entries in the secondary index database. This parameter is optional.

EXTRTN=

This parameter (not used in the example) specifies a user-exit routine. The user routine gets control after a source segment is built. The routine is used to suppress entries in the secondary index database when you cannot use the values that can be specified in the NULLVAL= parameter. This parameter is optional.

In the example, shown in Figure 131, a system-related field (/SX1) is used on the SUBSEQ parameter. System-related fields must also be coded on FIELD statements after the SEGM for the source segment. For more details, see "Making Keys Unique Using System Related Fields" on page 202.

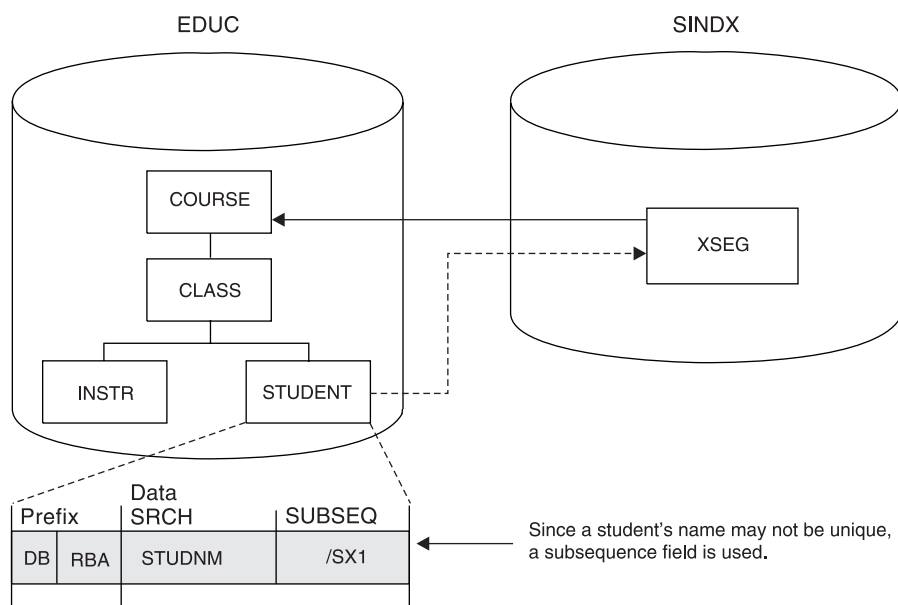


Figure 131. Databases for Secondary Indexing Example

Figure 132 shows the EDUC DBD for the example in Figure 131 on page 212.

```

DBD    NAME=EDUC,ACCESS=HDAM,...
SEGM   NAME=COURSE,...
FIELD  NAME=(COURSECD,...
LCHILD NAME=(XSE,SINDX),PTR=INDX
XDFLD  NAME=XSTUDENT,SEGMENT=STUDENT,SRCH=STUDNM,SUBSEQ=/SX1
SEGM   NAME=CLASS,...
FIELD  NAME=(EDCTR,...
SEGM   NAME=INSTR,...
FIELD  NAME=(INSTNO,...
SEGM   NAME=STUDENT,...
FIELD  NAME=SEQ,...
FIELD  NAME=STUDNM,BYTES=20,START=1
FIELD  NAME=/SX1

DBDGEN
FINISH
END

```

Figure 132. EDUC DBD for Secondary Indexing

Figure 133 shows the SINDX DBD for the example in Figure 131 on page 212.

```

DBD    NAME=SINDX,ACCESS=INDEX
SEGM   NAME=XSEG,...
FIELD  NAME=(XSEG,SEQ,U),BYTES=24,START=1
LCHILD NAME=(COURSE,EDUC),INDEX=XSTUDNT,PTR=SNGL

DBDGEN
FINISH
END

```

Figure 133. SINDX DBD for Secondary Indexing

DBD for the SINDX Database

- **DBD statement.** The DBD statement specifies the name of the secondary index database in the NAME= parameter. The ACCESS= parameter is always ACCESS=INDEX for the secondary index DBD.
- **SEGM statement.** You choose what is used in the NAME= parameter. This value is used when processing the secondary index as a separate database.
- **FIELD statement.** The NAME= parameter specifies the sequence field of the secondary index. In this case, the sequence field is composed of both the search and subsequence field data, the student name, and the system-related field /SX1. You specify what is chosen by NAME=parameter.
- **LCHILD statement.** The LCHILD statement specifies the name of the target, SEGM, and the name of the target database in the NAME= parameter. The INDEX= parameter has the name on the XDFLD statement in the target database. If the pointer segment contains a direct-address pointer to the target segment, the PTR= parameter is PTR=SNGL. The PTR= parameter is PTR=SYMB if the pointer segment contains a symbolic pointer to the target segment.

Choosing Secondary Indexes Versus Logical Relationships

While learning about secondary indexes and logical relationships, you might have noted that both options give you logical data structures. A logical data structure is a hierarchic data structure different from the data structure represented by the physical DBD. How, then, do you decide when to use a logical relationship and when to use a secondary index? This decision is based primarily on how your applications need to process the data.

When to Use a Secondary Index

In analyzing application requirements, if more than one candidate exists for the sequence field of a segment, use a secondary index. Choose one sequence field to be defined in the physical DBD. Then set up a secondary index to allow processing of the same segment in another sequence. For the example shown in Figure 134, access the customer segment that follows in both customer number (CUSTNO) and customer name (CUSTNAME) sequence. To do this, define CUSTNO as the sequence field in the physical DBD and then define a secondary index that processes CUSTOMER segments in CUSTNAME sequence.

| | | |
|----------|----------|--|
| CUSTOMER | | |
| CUSTNO | CUSTNAME | |

Figure 134. Fields in the CUSTOMER Segment

When to Use a Logical Relationship

If you have applications such as a bill-of-materials using a recursive structure, use a logical relationship. A recursive structure exists when there is a many-to-many association between two segments in the same physical hierarchy. For example, in the segments shown in Figure 135, the assembly “car” is composed of many parts, one of which is an engine. However, the engine is itself an assembly composed of many parts.




Figure 135. Assembly and Parts as Examples to Demonstrate Segments Logical Relationship

Related Reading: Recursive structure are explained in “Recursive Structures: Same Database Logical Relationships” on page 172.

Finally, you can have application requirements that result in a segment that appears to have two parents. In the example shown in Figure 136 on page 215, the customer database keeps track of orders (CUSTORDN). Each order can have one or more line items (ORDLINE), with each line item specifying one product (PROD) and model (MODEL). In the product database, many outstanding line item requests can exist for a given model. This type of relationship is called a many-to-many relationship and is handled in IMS through a logical relationship.

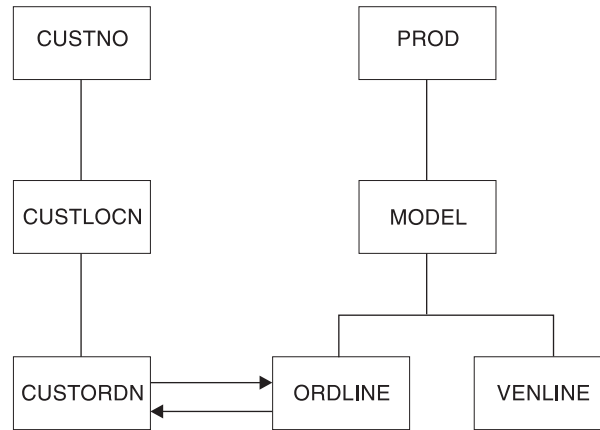


Figure 136. Example of a Segment That Appears to Have Two Parents

Variable-Length Segments

Database types that support variable-length segments:

- HISAM
- SHISAM
- HDAM
- PHIDAM
- HIDAM
- PHDAM
- DEDB

Variable-length segments are simply segments whose length can vary in occurrence of some segment types. A database can contain both variable-length segment and fixed-length segment types. Variable-length segments can be used for HISAM, HDAM, PHDAM, HIDAM, and PHIDAM databases.

How to Specify Variable-Length Segments

It is the data portion of a variable-length segment whose length varies. The data portion varies between a minimum and a maximum number of bytes. As shown in Figure 137, you specify minimum and maximum size in the BYTES= keyword in the SEGM statement in the DBD. Because IMS needs to know the length of the data portion of a variable-length segment, you include a 2-byte size field in each segment when loading it. The size field is in the data portion of the segment. The length of the data portion you specify must include the two bytes used for the size field. If the segment type has a sequence field, the minimum length specified in the size field must equal at least the size field and all data to the end of the sequence field.




Figure 137. How Variable-Length Segments Are Specified

How Variable-Length Segments Are Stored and Processed

When a variable-length segment is initially loaded, the space used to store its data portion is the length specified in the MINBYTES operand or the length specified in the size field, whichever is larger. If the space in the MINBYTES operand is larger, more space is allocated for the segment than is required. The additional space can be used when existing data in the segment is replaced with data that is longer.

The prefix and data portion of HDAM, PHDAM, HIDAM, and PHIDAM variable-length segments can be separated in storage when updates occur. When this happens, the first four bytes following the prefix point to the separated data portion of the segment.

Figure 138 shows the format of a HISAM variable-length segment. It is also the format of an HDAM, PHDAM, HIDAM, or PHIDAM variable-length segment when the prefix and data portion of the segment have not been separated in storage.




Figure 138. Format of HISAM Variable-Length Segments

Figure 139 shows the format of an HDAM, PHDAM, HIDAM, or PHIDAM variable-length segment when the prefix and data portion of the segment have been separated in storage.

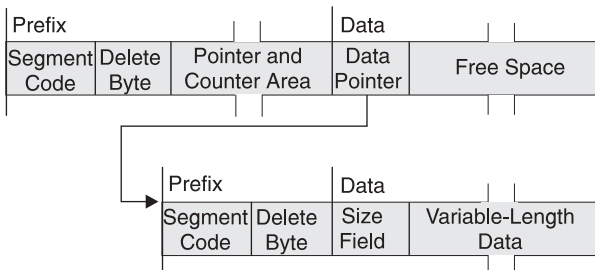


Figure 139. Format of HDAM, PHDAM, HIDAM or PHIDAM Variable-Length Segments

After a variable-length segment is loaded, replace operations can cause the size of data in it to be either increased or decreased. When the length of data in an existing HISAM segment is increased, the logical record containing the segment is rewritten to acquire the additional space. Any segments displaced by the rewrite are put in overflow storage. Displacement of segments to overflow storage can affect performance. When the length of data in an existing HISAM segment is decreased, the logical record is rewritten so all segments in it are physically adjacent.

When a replace operation causes the length of data in an existing HDAM, PHDAM, HIDAM, or PHIDAM segment to be increased, one of two things can happen:

- If the space allocated for the existing segment is long enough for the new data, the new data is simply placed in the segment. This is true regardless of whether the prefix and data portions of the segment were previously separated in the data set.
- If the space allocated for the existing segment is not long enough for the new data, the prefix and data portions of the segment are separated in storage. IMS

puts the data portion of the segment as close to the prefix as possible. Once the segment is separated, a pointer is placed in the first four bytes following the prefix to point to the data portion of the segment. This separation increases the amount of space needed for the segment, because, in addition to the pointer kept with the prefix, a 1-byte segment code and 1-byte delete code are added to the data portion of the segment (see Figure 138 on page 216). In addition, if separation of the segment causes its two parts to be stored in different blocks, two read operations will be required to access the segment.

When a replace operation causes the length of data in an existing HDAM, PHDAM, HIDAM, or PHIDAM segment to be decreased, one of three things can happen:

- If prefix and data are not separated, the data in the existing segment is replaced with the new, shorter data followed by free space.
- If prefix and data are separated but sufficient space is not available immediately following the original prefix to recombine the segment, the data in the separated data portion of the segment is replaced with the new, shorter data followed by free space.
- If prefix and data are separated and sufficient space is available immediately following the original prefix to recombine the segment, the new data is placed in the original space, overlaying the data pointer. The old separated data portion of the segment is then available as free space in HD databases.

When to Use Variable-Length Segments

Use variable-length segments when the length of data in your segment varies, for example, with descriptive data. By using variable-length segments, you do not need to make the data portion of your segment type as long as the longest piece of descriptive data you have. This saves storage space. Note, however, that if you are using HDAM, PHDAM, HIDAM, or PHIDAM databases and your segment data characteristically grows in size over time, segments will split. If a segment split causes the two parts of a segment to be put in different blocks, two read operations will be required to access the segment until the database is reorganized. So variable-length segments work well if segment size varies but is stable (as in an address segment). Variable-length segments might not work well if segment size typically grows (as in a segment type containing a cumulative list of sales commissions).

What Application Programmers Need to Know about Variable-Length Segments

If you are using variable-length segments in your database, you need to let application programmers who will be using the database know this. They need to know which of the segment types they have access to are variable in length and the maximum size of each of these variable-length segment types. In calculating the size of their I/O area, application programmers must use the maximum size of a variable-length segment. In addition, they need to know that the first two bytes of the data portion of a variable-length segment contain the length of the data portion including the size field.

Working with the application programmer, you should devise a scheme for accessing data in variable-length segments. You should devise a scheme because if variable-length fields and fixed-length fields in a segment are mixed, the application program has no way of knowing where specific fields begin. One way to solve this problem is to put the size of a variable-length field at the beginning of the variable-length field. If a segment has only one variable-length field, it can be made

Variable-Length Segments

the last field in the segment. If it is at all possible, the simplest scheme is to have only one field in a variable-length segment.

Adding or Converting to Variable-Length Segments

Information on how to add variable-length segments to an existing database and convert an entire database to variable-length segments is in Chapter 16, “Modifying Databases,” on page 429.

Segment Edit/Compression Exit Routine

The following database types support the Segment Edit/Compression exit routine:

- HISAM
- HDAM
- PHDAM
- HIDAM
- PHIDAM
- DEDB

Detailed information on how the Segment Edit/Compression exit routine works and how you use it is in *IMS Version 9: Customization Guide*. This topic introduces you to the facility.

The Segment Edit/Compression exit routine allows you to encode, edit, or compress the data portion of a segment. You can use this facility on segment data in full function databases and Fast Path DEDBs. You write the routine (your edit routine) that actually manipulates the data in the segment. The IMS code gives your edit routine information about the segment’s location and assists in moving the segment back and forth between the buffer pool and the application program’s I/O area.

The Segment Edit/Compression exit routine lets you:

- **Encode data for security purposes.** Encoding data consists of “scrambling” segment data when it is on the device so only programs with access to the edit routine can see it in decoded form.
- **Edit data.** Editing data allows application programs to receive data in a format other than the one in which it is stored. For example, an application program might receive segment fields in an order other than the one in which they are stored; an application program might require all blank space be removed from descriptive data.
- **Compress data.** This allows better use of DASD storage because segments can be compressed when written to the device and then expanded when passed back to the application program. Segment data might be compressed, for example, by removing all blanks and zeros.
- **Expand Data.** The DEDB Sequential Dependent Scan utility invokes the Segment Edit/Compression exit routine (DFSCMPX0) to expand compressed SDEP segments when you specify both SDEP segment compression in the DBD and the DEDB Scan utility keyword, EXPANDSEG.

Related Reading: EXPANDSEG and the DEDB Scan utility are described in *IMS Version 9: Utilities Reference: Database and Transaction Manager*. The segment compression exit is described in *IMS Version 9: Customization Guide*.

Two types of segment manipulation are possible using the Segment Edit/Compression exit routine:

- **Data compression**— movement or compression of data within a segment in a manner that does not alter the content or position of the key field. Typically, this involves compression of data from the end of the key field to the end of the segment. When a fixed-length segment is compressed, a 2-byte field must be added to the beginning of the data portion of the segment by the user data compression routine. This field is used by IMS to determine secondary storage requirements and is the only time that the location of the key field can be altered. The segment size field of a variable-length segment cannot be compressed but must be updated to reflect the length of the compressed segment.
- **Key compression**— movement or compression of any data within a segment in a manner that can change the relative position, value, or length of the key field and any other fields except the size field. The segment size field of a variable-length segment must be updated by the compression routine to reflect the length of the compressed segment.

You specify the use of the segment edit/compression exit routine when you define a segment type. Any segment type can be edited or compressed (using either data or key compression) as long as the segment is:

- Not a logical child
- Not in an HSAM, SHISAM, or index database

The use of the segment edit/compression exit routine is defined in physical database DBDs. This exit routine's use cannot be defined in a logical database DBD.

Data compression is allowed but key compression is not allowed when the segment is:

- A root segment in a HISAM database
- A segment in a DEDB database

Things to Consider Before Using the Segment Edit/Compression Exit Routine

Before using a Segment Edit/Compression routine you should be aware of the following points:

- Because your edit routine is executed as part of a DL/I call, if it abnormally terminates so does the entire IMS region.
- Your routine cannot use the operating system macros LOAD, GETMAIN, SPIE or STAE.

Related Reading: For alternatives to these macros, see *IMS Version 9: Customization Guide*.

- Editing and compressing of each segment on its way to or from an application program requires additional processor time.

Depending on the options you select, search time to locate a specific segment can increase. If you are fully compressing the segment using key compression, every segment type that is a candidate to satisfy either a fully qualified key or data field request must be expanded or divided. IMS then examines the appropriate field. For key field qualification, only those fields from the start of the segment through the sequence field are expanded during the search. For data field qualification, the total segment is expanded. In the case of data compression and a key field request, little more processing is required to locate the segment than that of non-compressed segments. Only the segment sequence field is used to determine if this segment occurrence satisfies the qualification.

Segment Edit/Compression Exit Routine

Other considerations can affect total system performance, especially in an online environment. For example, being able to load an algorithm table into storage gives the compression routine a large amount of flexibility. However, this can place the entire IMS control region into a wait state until the requested member is present in storage. It is suggested that all alternatives be explored to lessen the impact of situations such as this.

Preventing Split Segments from Impacting Performance

When segments in a full-function database grow larger than the size of their current location, replace calls can split the segments. When segments are split, their prefixes remain in their existing location, but their data parts are stored in a new location, possibly in another block or CI. Split segments can negatively affect performance by requiring additional reads to retrieve both parts of the segments.

To prevent IMS from splitting compressed segments, you can specify a minimum size for the segments that includes extra padded space. This gives the compressed segment room to grow and decreases the chance that IMS will split the segment.

You specify the minimum size for fixed-length full-function segments differently than you do for variable-length full-function segments:

- For fixed-length segments, specify the minimum size using both the fourth and fifth subparameters on the COMPRTN= parameter of the SEGM statement. The fourth subparameter, *size*, only defines the minimum size if you also specify the fifth subparameter, PAD.
- For variable-length segments, specify the minimum size using the second subparameter, *min_bytes*, of the BYTES= parameter of the SEGM statement.

Related Reading: For a complete description of the COMPRTN= and BYTES= parameters of SEGM statements for full-function databases, see *IMS Version 9: Utilities Reference: System*.

DEDB segments are never split by replace calls. If a DEDB segment grows beyond the size of its current location, the entire segment, including its prefix, is moved to a new location. For this reason, it is not necessary to pad compressed DEDB segments.

How to Specify the Segment Edit/Compression Exit Routine

To specify the use of the segment edit/compression exit routine for a segment, use the COMPRTN= keyword of the SEGM statement in the DBD.

Related Reading: For more information on using the COMPRTN= keyword to specify the use of the segment edit/compression exit routine, see *IMS Version 9: Utilities Reference: System*.

Converting to the Segment Edit/Compression Exit Routine

Information on how to convert an existing database so it can use a Segment Edit/Compression exit routine is discussed in Chapter 16, “Modifying Databases,” on page 429.

Data Capture Exit Routines

The following database types support data capture exit routines:

- HISAM
- SHISAM
- HDAM
- PHDAM
- HIDAM
- PHIDAM
- DEDB

The Data Capture exit routine is an installation-written exit routine. Data Capture exit routines promote and enhance database coexistence. Data Capture exit routines capture segment-level data from a DL/I database for propagation to DB2 UDB for z/OS databases. Installations running IMS and DB2 UDB for z/OS databases can use Data Capture exit routines to exchange data across the two database types.

Data Capture exit routines can be written in assembler language, C, VS COBOL II, or PL/I. *IMS Version 9: Customization Guide* describes Data Capture exit routines in detail.

Data Capture exit routines are supported by IMS Transaction Manager and Database Manager. DBCTL support is for BMPs only.

Data Capture exit routines are compatible with the following physical database structures:

HIDAM
PHIDAM
HISAM
SHISAM
DEDB

Data Capture exit routines do not support segments in secondary indexes.

A Data Capture exit routine is called based on segment-level specifications in the DBD. When a Data Capture exit routine is specified on a database segment, it is invoked by all application program activity on that segment, regardless of which PSB is active. Therefore, Data Capture exit routines are global. Using a Data Capture exit routine can have a performance impact across the entire database system.

DBD Parameters for Data Capture Exit Routines

This topic contains programming interface information.

Using Data Capture exit routines requires specification of one or two DBD parameters and subsequent DBDGEN. The EXIT= parameter identifies which Data Capture exit routines will run against segments in a database. The VERSION= parameter records important information about the DBD for use by Data Capture exit routines.

The EXIT= Parameter

To use the Data Capture exit routine, you must use the optional EXIT= parameter. You specify EXIT= on either the DBD or SEGM statements of physical database definitions.

Specifying EXIT= on the DBD statement applies a Data Capture exit routine to all segments within a database structure. Specifying EXIT= on the SEGM statement applies a Data Capture exit routine to only that segment type.

You can override Data Capture exit routines specified on the DBD statement by specifying EXIT= on a SEGM statement. EXIT=NONE on a SEGM statement cancels all Data Capture exit routines specified on the DBD statement for that segment type. A physical child does not inherit an EXIT= parameter specified on the SEGM statement of its physical parent.

You can specify multiple Data Capture exit routines on a single DBD or SEGM statement. For example, you might code a DBD statement as:

```
DBD   EXIT=((EXIT1A),(EXIT1B))
```

The name of the Data Capture exit routine that you intend to use is the only required operand for the EXIT= parameter. Exit names can have a maximum of eight alphanumeric characters. For example, if you specify a Data Capture exit routine with the name EXITA on a SEGM statement in a database, the EXIT= parameter is coded as follows:

```
SEGM   EXIT=(EXITA,KEY,DATA,NOPATH,(CASCADE,KEY,DATA,NOPATH))
```

KEY, NOPATH, DATA, CASCADE, KEY, DATA, and NOPATH are default operands. These defaults define what data is captured by the exit routine when a segment is updated by an application program.

Related Reading:

- For more information about the Data Capture exit routine, see *IMS Version 9: Customization Guide*.
- For a full description of the EXIT= parameter on both the DBD and SEGM statements, see *IMS Version 9: Utilities Reference: System*.

The VERSION= Parameter

VERSION= is an optional parameter that supports Data Capture exit routines. VERSION= is specified on the DBD statement as:

```
VERSION='character string'
```

The maximum length of the character string is 255 bytes. You can use VERSION= to create a naming convention that denotes the database characteristics that affect the proper functioning of Data Capture exit routines. You might use VERSION= to flag DBDs containing logical relationships, or to indicate which data capture exit routines are defined on the DBD or SEGM statements. VERSION= might be coded as:

```
DBD   VERSION='DAL-&SYSDATE-&SYSTIME'
```

DAL, in this statement, tells you that Data Capture exit routine A is specified on the DBD statement (D), and that the database contains logical relationships (L). &SYSDATE and &SYSTIME tell you the date and time the DBD was generated.

If you do not specify a VERSION= parameter, DBDGEN generates a default 13-character date-time stamp. The default consists of an 8-byte date stamp and a 5-byte time stamp with the following format:

MM/DD/YYHH.MM

The default date-time stamp on VERSION= is identical to the DBDGEN date-time stamp.

VERSION= is passed as a variable length character string with a 2-byte length of the VERSION=, which does not include the length of the LL.

Call Sequence of Data Capture Exit Routines

This topic contains programming interface information.

A Data Capture exit routine is invoked once per segment update for each segment for which the Data Capture exit routine is specified. Data Capture exit routines are invoked multiple times for a single call under certain conditions. These conditions include:

- Path updates.
- Cascade deletes when multiple segment types or multiple segment occurrences are deleted.
- Updates on logical children.
- Updates on logical parents.
- Updates on a single segment when multiple Data Capture exit routines are specified against that segment. Each exit is invoked once, in the order it is listed on the DBD or SEGM statements.

When multiple segments are updated in a single application program call, Data Capture exit routines are invoked in the same order in which IMS physically updates the segments:

1. Path inserts are executed “top-down” in DL/I. Therefore, a Data Capture exit routine for a parent segment is called before a Data Capture exit routine for that parent’s dependent.
2. Cascade deletes are executed “bottom-up”. All dependent segments’ exits are called before their respective parents’ exits on cascade deletes. IMS physically deletes dependent segments on cascade deletes only after it has validated the delete rules by following the hierarchy to the lowest level segment. After delete rules are validated, IMS deletes segments starting with the lowest level segment in a dependent chain and continuing up the chain, deleting the highest level parent segment in the hierarchy last. Data Capture exit routines specified for segments in a cascade delete are called in reverse hierarchical order.
3. Path replaces are performed “top-down” in IMS. In Data Capture exit routines defined against segments in path replaces, parent segments are replaced first. All of their descendents are then replaced in descending hierarchical order.

When an application program does a cascade delete on logically related segments, Data Capture exit routines defined on the logical child are always called before Data Capture exit routines defined on the logical parent. Data Capture exit routines are called even if the logical child is higher in the physical hierarchy, except in recursive structures where the delete results in the deletion of a parent of the deleted segment.

Data Passed To and Captured By the Data Capture Exit Routine

This topic contains programming interface information.

Data is passed to Data Capture exit routines when an application program updates IMS with a DL/I insert, delete, or replace call. Segment data passed to Data Capture exit routines is always physical data. When the update involves logical children, the data passed is physical data and the concatenated key of the logical parent segment. For segments that use the Segment Edit/Compression exit routine (DFSCMPX0), the data passed is expanded data.

When an application replaces a segment, both the existing and the replacement physical data are captured. In general, segment data is captured even if the application call does not change the data. However, for full-function databases, IMS compares the before and after data. If the data has not changed, IMS does not update the database or log the replace data. Because data is not replaced, Data Capture exit routines specified for that segment are not called and the data is not captured.

Data might be captured during replaces even if segment data does not change when:

1. The application inserts a concatenation of a logical child and logical parent, IMS replaces the logical parent, and the parent data does not change.
2. The application issues a replace for a segment in a DEDB database.

In each case, IMS updates the database without comparing the before and after data, and therefore the data is captured even though it does not change.

The entire segment, before and after, is passed to Data Capture exit routines when the application replaces a segment. When the exit routine is interested in only a few fields, it is recommended that the SQL update request not be issued until after the before and after replace data for those fields is compared to see if the fields were changed.

Data Capture Call Functions

This topic contains programming interface information.

Data Capture exit routines are called when segment data is updated by an application program insert, replace, or delete call. Optionally, Data Capture exit routines are called when DL/I deletes a dependent segment because the application program deleted its parent segment, a process known as cascade delete. Data Capture exit routines are passed two functions to identify the following:

1. The action performed by the application program
2. The action performed by IMS

The two functions that are passed to the Data Capture exit routines are:

- Call function. The DL/I call, ISRT, REPL, or DLET, that is issued by the application program for the segment.
- Physical function. The physical action, ISRT, REPL, or DLET, performed by IMS as a result of the call. The physical function is used to determine the type of SQL request to issue when propagating data.

The call and physical functions passed to the exit routine are always the same for replace calls. However, the functions passed might differ for delete or insert calls:

- For delete calls resulting in cascade deletes, the call function passed is CASC (to indicate the cascade delete) and the physical function passed is DLET.
- For insert calls resulting in the insert of a logical child and the replace of a logical parent (because the logical parent already exists), the call function passed is ISRT and the physical function passed is REPL. IMS physically replaces the logical parent with data inserted by the application program even if the parent data does not change. Both call and physical functions are then used, based on the data propagation requirements, to determine the SQL request to issue in the Data Capture exit routine.

Cascade Delete When Crossing Logical Relationships

This topic contains programming interface information.

If the EXIT= options specify NOCASCADE, data is not captured for cascade deletes. However, when a cascade delete crosses a logical relationship into another physical database to delete dependent segments, a Data Capture exit routine needs to be called in order to issue the SQL delete for the parent of the physical structure in DB2 UDB for z/OS. Rather than requiring the EXIT= CASCADE option, IMS always calls the exit routine for a segment when deleting the parent segment in a physical database record with an exit routine defined, regardless of the CASCADE/NOCASCADE option specified on the segment. IMS bypasses the NOCASCADE option only when crossing logical relationships into another physical database. As with all cascade deletes, the call function passed is CASC and the physical function passed is DLET.

Data Capture Exit Routines and Logically Related Databases

This topic contains programming interface information.

Segment data passed to Data Capture exit routines is always physical data. Consequently, you must place restrictions on delete rules in logically related databases supporting Data Capture exit routines. Table 17 summarizes which delete rules you can and cannot use in logically related databases with Data Capture exit routines specified on their segments.

Table 17. Delete Rule Restrictions for Logically Related Databases Using Data Capture Exit Routines

| Segment Type | Virtual Delete Rule | Logical Delete Rule | Physical Delete Rule |
|------------------|---------------------|---------------------|----------------------|
| Logical Children | Yes | No | No |
| Logical Parents | No | Yes | Yes |

When a logically related database has a delete rule violation on a logical child:

- The logical child cannot have a Data Capture exit routine specified.
- No ancestor of the logical child can have a Data Capture exit routine specified.

When a logically related database has a delete rule violation on a logical parent, the logical parent cannot have a Data Capture exit routine specified. ACBGEN validates logical delete rule restrictions and will not allow a PSB that refers to a database that violates these restrictions to proceed.

Converting to Data Capture Exit Routines

Related Reading:

Data Capture Exit Routines

- For information on how to convert an existing database for Data Capture exit routines, see “Converting Databases for Data Capture Exit Routines and Asynchronous Data Capture” on page 453.
- For detailed information on coding the EXIT= and VERSION= parameters, see *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

Field-Level Sensitivity

The following database types support field-level sensitivity:

- HSAM
- HISAM
- SHISAM
- HDAM
- PHDAM
- HIDAM
- PHIDAM

Field-level sensitivity gives you an increased level of data independence by isolating application programs from:

- Changes in the arrangement of fields within a segment
- Addition or deletion of data within a segment

In addition, field-level sensitivity enhances data security by limiting an application program to a subset of fields within a segment, and controlling replace operations at the field level.

Field-level sensitivity allows you to reformat a segment type. Reformatting a segment type can be done without changing the application program's view of the segment data, provided fields have not been removed or altered in length or data type. Fields can be added to or shifted within a segment in a manner transparent to the application program. Field-level sensitivity gives applications a segment organization that always conforms to what is specified in the SENFLD statements. (SENFLD statements are described in “How to Specify Use of Field-Level Sensitivity in the DBD and PSB” on page 227, but basically they determine the order of fields in a segment as seen by an application program.)

Using Field-Level Sensitivity as a Mapping Interface

Field-level sensitivity acts as a mapping interface by letting PSBGEN field locations differ from DBDGEN field locations. Mapping is invoked after the segment edit routine on input and before the segment edit routine on output. When creating a sequential data set from database information (or creating database information from a sequential data set), field-level sensitivity can reduce or eliminate the amount of formatting an application program must do.

Using Field-Level Sensitivity with Variable-Length Segments

If field-level sensitivity is used with variable-length segments, you can add new fields to a segment without reorganizing the database. FIELD definitions in a DBDGEN allow you to enlarge segment types without affecting any previous users of the segment. The DBDGEN FIELD statement lets you specify a field that doesn't yet exist in the physical segment but that will be dynamically created when the segment is retrieved.

Field-level sensitivity can help in the transition of an application program from a non-database environment to a database environment. Application programs that formerly accessed z/OS files might be able to receive the same information in the same format if the database was designed with conversion in mind.

Field-level sensitivity is *not* supported for DEDBs and MSDBs.

How to Specify Use of Field-Level Sensitivity in the DBD and PSB

An application program's view of data is defined through the PSBGEN utility using SENFLD statements following the SENSEG statement. In the SENFLD statement, the NAME= parameter identifies a field that has been defined in the segment through the DBDGEN utility.

The START= parameter defines the starting location of the field in the application program's I/O area. In the I/O area, fields do not need to be located in any particular order, nor must they be contiguous. The end of the segment in the I/O area is defined by the end of the right most field. All segments using field-level sensitivity appear fixed in length in the I/O area. The length is determined by the sum of the lengths of fields on SENFLD statements associated with a SENSEG statement.

Figure 140 is an example of field-level sensitivity. Following the figure is information about coding field-level sensitivity.

Field-level sensitivity is used below to reposition three fields from a physical segment in the application program's I/O area.




Figure 140. DBD and PSB Coding for Field-Level Sensitivity

Figure 141 shows the DBD for the example shown in Figure 140.

```
SEGM  NAME=EMPREC,BYTES=100
FIELD NAME=(EMPNO,SEQ),BYTES=5,START=1,TYPE=C
FIELD NAME=EMPNAME,BYTES=20,START=6,TYPE=C
FIELD NAME=BIRTHD,BYTES=6,START=26,TYPE=C
FIELD NAME=SAL,BYTES=3,START=32,TYPE=P
FIELD NAME=ADDRESS,BYTES=60,START=41,TYPE=C
```

Figure 141. DBD Example for Field-Level Sensitivity

Figure 142 shows the PSB for the figure shown in Figure 140.

```
SENSEG NAME=EMPREC,PROCOPT=A
SENFLD NAME=EMPNAME,START=1,REPL=N
SENFLD NAME=EMPNO,START=25
SENFLD NAME=ADDRESS,START=35,REPL=Y
```

Figure 142. PSB Example for Field-Level Sensitivity

Field-Level Sensitivity

- A SENFLD statement is coded for each field that can appear in the I/O area. A maximum of 255 SENFLD statements can be coded for each SENSEG statement, with a limit of 10000 SENFLD statements for a single PSB.
- The optional REPL= parameter on the SENFLD statement indicates whether replace operations are allowed on the field. In the figure, replace is not allowed for EMPNAME but is allowed for EMPNO and ADDRESS. If REPL= is not coded on a SENFLD statement, the default is REPL=Y.
- The TYPE= parameter on FIELD statements in the DBD is used to determine fill values on insert operations.

Retrieving Segments Using Field-Level Sensitivity

When you retrieve segments using field-level sensitivity, you should be aware of the following information:

- Gaps between fields in the I/O area are set to blanks on a retrieve call.
- If an application program uses a field in an SSA, that field must be coded on a SENFLD statement. This rule does not apply to sequence fields used in an SSA on retrieve operations.

Figure 143 shows an example of a retrieve call based on the DBD and PSB in Figure 140.

| Physical Segment on Disk | | | | | |
|--------------------------|-----------------------|------------------|------------|-------|---------------------|
| EMPNO 12345 | EMPNAME SMITH, JOE | BIRTHD 480207 | SAL 999 | | ADDRESS NEW YORK |
| 1 | 5 6 | 25 | 26 31 | 32 34 | 41 100 |
| Bytes | | | | | |

| I/O Area After Get Call | | | | | |
|-------------------------|-------|----------------|-------|----|---------------------|
| EMPNAME SMITH, JOE | ↔ ↔ ↔ | EMPNO 12345 | ↔ ↔ ↔ | | ADDRESS NEW YORK |
| 1 | 20 | 25 29 | | 35 | 94 |
| Bytes | | | | | |

Figure 143. Example of a Retrieve Call

Replacing Segments Using Field-Level Sensitivity

The SENFLD statement must allow replace operations (REPL=Y) if the application program is going to replace data in a segment. In Figure 140 on page 227, the SENFLD statement for EMPNAME specifies REPL=N. A “DA” status code would be returned if the application program tried to replace the EMPNAME field. Figure 144 on page 229 shows an example of a REPL call based on the DBD and PSB in Figure 140.

| Physical Segment on Disk | | | | | |
|--------------------------|------------|--------|-------|-------|----------|
| EMPNO | EMPNAME | BIRTHD | SAL | | ADDRESS |
| 12345 | SMITH, JOE | 480207 | 999 | | NEW YORK |
| 1 | 5 6 | 25 | 26 31 | 32 34 | 41 100 |
| Bytes | | | | | |

| I/O Area | | | | | |
|------------|------|-------|------|----|----------|
| EMPNAME | | EMPNO | | | ADDRESS |
| SMITH, JOE | ↔↔↔↔ | 12345 | ↔↔↔↔ | | NEW YORK |
| 1 | 20 | 25 29 | | 35 | 94 |
| Bytes | | | | | |

| Physical Segment on Disk After Update | | | | | |
|---------------------------------------|------------|--------|-------|-------|----------|
| EMPNO | EMPNAME | BIRTHD | SAL | | ADDRESS |
| 12345 | SMITH, JOE | 480207 | 999 | | NEW YORK |
| 1 | 5 6 | 25 | 26 31 | 32 34 | 41 100 |
| Bytes | | | | | |

Figure 144. Example of a REPL Call

Inserting Segments Using Field-Level Sensitivity

The TYPE= parameter on the SEGM statement of the DBD determines the fill value in the physical segment when an application program is not sensitive to a field on insert calls.

| TYPE= | Fill Value |
|-------|---------------------|
| X | Binary Zeros |
| P | Packed Decimal Zero |
| C | Blanks |

The fill value in the physical segment is binary zeros when:

- Space in a segment is not defined by a FIELD macro in the DBD
- A defined DBD field is not referenced on the insert operation

Figure 145 on page 230 shows an example of an insert operation based on the DBD and PCB in Figure 140 on page 227.

Field-Level Sensitivity

| I/O Area | | | | | |
|------------------------|----|----------------|----|--------------------|----|
| EMPNAME ADAMS, DICK | | EMPNO 23456 | | ADDRESS VERMONT | |
| 1 | 20 | 25 | 29 | 35 | 94 |
| Bytes | | | | | |

| Physical Segment on Disk After Update | | | | | |
|---------------------------------------|------------------------|---------------|------------|-------|--------------------|
| EMPNO 23456 | EMPNAME ADAMS, DICK | BIRTHD X↔X | SAL 00+ | 0↔0 | ADDRESS VERMONT |
| 1 | 5 6 | 25 | 26 31 | 32 34 | 41 100 |
| Bytes | | | | | |

Figure 145. Example of an ISRT Call

Blanks are inserted in the BIRTHD field because its FIELD statement in the DBD specifies TYPE=C. Packed decimal zero is inserted in the SAL field because its FIELD statement in the DBD specifies TYPE=P. Binary zeros are inserted in positions 35 to 40 because no FIELD statement was coded for this space in the DBD.

Using Field-Level Sensitivity When Fields Overlap

On the SENFLD statement, you code the starting position of fields as they will appear in the I/O area. If fields overlap in the I/O area, here are the rules you must follow:

- Two different bytes of data cannot be moved to the same position in the I/O area on input.
- The same data can be moved to different positions in the I/O area on retrieve operations.
- Two bytes from different positions in the I/O area cannot be moved to the same DBD field on output.

Using Field-Level Sensitivity When Path Calls Are Issued

If an application program issues path calls while using field level sensitivity, here are the rules you must follow:

- You should not code SENFLD statements so that two fields from different physical segments are in the same segment in the I/O area.
- PROCOPT=P is required on the PCB statement.

Using Field-Level Sensitivity with Logical Relationships

Here are the rules you must follow when using field-level sensitivity with segments involved in a logical relationship:

- Application programs can not insert sensitive to a logical child.
- The same field can be referenced in more than one SENFLD statement within a SENSEG. If the duplicate field names are part of a concatenated segment and the same field name appears in both parts of the concatenation, the first part references the logical child. The second and all subsequent parts reference the logical parent. This referencing sequence determines the order in which fields are moved to the I/O area.
- When using field-level sensitivity with a virtual logical child, the field list of the paired segment is searched after the field list of the virtual segment and before the field list of the logical parent.

Using Field-Level Sensitivity with Variable-Length Segments

When field-level sensitivity is used with a variable-length segment, an application program's view of the segment is fixed in length and does not include the 2-byte length field. This topic and its subtopics address special situations when field level sensitivity is used with variable-length segments. First, however, here is some general information about using field-level sensitivity with variable-length segments:

- When inserting a variable-length segment, the length used is the minimum length needed to hold all sensitive fields.
- When replacing a variable-length segment, if the length has to be increased to contain data an application program has modified, the length used is the minimum length needed to hold the modified data.
- An application program cannot be sensitive to overlapping fields in a variable-length segment with get or update sensitivity if the data type of any of those fields is not character.
- Existing programs processing variable-length segments that use the length field to determine the presence or absence of a field might need to be modified if segments are inserted or updated by programs using field-level sensitivity.

When field-level sensitivity is used with variable-length segments, two situations exist that you should know about. The first is when fields are missing. The second is when fields are partially present. This topic examines the following information:

- Retrieving Missing Fields
- Replacing Missing Fields
- Inserting Missing Fields
- Retrieving Partially Present Fields
- Replacing Partially Present Fields

Retrieving Missing Fields

If a field does not exist in the physical variable-length segment at retrieval time, the corresponding field in the application program's I/O area is filled with a value based on the data type specified in the DBD. Figure 146 is an example of a missing field on a retrieve call based on the DBD and PSB in Figure 147 and Figure 148 on page 232.

| Physical Segment on Disk | | | | |
|--------------------------|-------|------------|---|---|
| LL | EMPNO | EMPNAME | | |
| 27 | 12345 | SMITH, JOE | | |
| 1 | 2 | 3 | 7 | 8 |
| Bytes | | | | |
| 27 | | | | |

| User I/O Area After Get Call | | | | |
|------------------------------|-------|----|---------|----|
| EMPNAME | EMPNO | | ADDRESS | |
| SMITH, JOE | 12345 | | ⌘↔⌘ | |
| 1 | 20 | 25 | 29 | 35 |
| Bytes | | | | |
| 94 | | | | |

Figure 146. Example of a Missing Field on a Retrieve Call

Field-Level Sensitivity

DBD

```
SEGM  NAME=EMPREC,BYTES=(102,7)
FIELD NAME=(EMPNO,SEQ),BYTES=5,START=3,TYPE=C
FIELD NAME=EMPNAME,BYTES=20,START=8,TYPE=C
FIELD NAME=BIRTHD,BYTES=6,START=28,TYPE=C
FIELD NAME=ADDRESS,BYTES=60,START=43,TYPE=C
```

Figure 147. DBD Example for Field-Level Sensitivity with Variable-Length Segments

PSB

```
SENSEG NAME=EMPREC,PROCOPT=A
SENFLD NAME=EMPNAME,START=1,REPL=N
SENFLD NAME=EMPNO,START=25
SENFLD NAME=ADDRESS,START=35,REPLY=Y
```

Figure 148. PSB Example for Field-Level Sensitivity with Variable-Length Segments

The length field is not present in the I/O area. Also, the address field is filled with blanks, because TYPE=C is specified on the FIELD statement in the DBD.

Replacing Missing Fields

A missing field that is not replaced does not affect the physical variable-length segment. Figure 149 is an example of a missing field on a replace call based on the DBD and PSB in Figure 147.

| Physical Segment on Disk Before Update | | | | |
|---|-------|------------|---|------|
| LL | EMPNO | EMPNAME | | |
| 27 | 12345 | SMITH, JOE | | |
| 1 | 2 | 3 | 7 | 8 27 |
| Bytes | | | | |

| User I/O Area | | | | |
|---------------|-------|-------|-------|---------|
| EMPNAME | | EMPNO | | ADDRESS |
| SMITH, JOE | ↔↔↔↔↔ | 12345 | ↔↔↔↔↔ | ↔↔↔↔↔ |
| 1 | 20 | 25 | 29 | 35 94 |
| Bytes | | | | |

| Physical Segment on Disk After Update | | | | |
|--|-------|------------|---|------|
| LL | EMPNO | EMPNAME | | |
| 27 | 12345 | SMITH, JOE | | |
| 1 | 3 | 3 | 7 | 8 27 |
| Bytes | | | | |

Figure 149. First Example of a Missing Field on a Replace Call

The length field, maintained by IMS, does not include room for the address field, because the field was missing and not replaced.

On a replace call, if a field returned to the application program with a fill value is changed to a non-fill value, the segment length is increased to the minimum size needed to hold the modified field.

- The 'LL' field is updated to include the full length of the added field and all fields up to the added field.
- The TYPE= parameter in the DBD (see Figure 147) determines the fill value for non-sensitive DBD fields up to the added field.

- Binary zero is the fill value for space up to the added field that is not defined by a FIELD statement in the DBD.

Figure 150 is an example of a missing field on a replace call based on the DBD and PSB in Figure 147 on page 232.

| Physical Segment on Disk Before Update | | | | | | | | | |
|--|---|----------------|---|-----------------------|--|--|--|--|----|
| LL 27 | | EMPNO 12345 | | EMPNAME SMITH, JOE | | | | | |
| 1 | 3 | 3 | 7 | 8 | | | | | 27 |
| Bytes | | | | | | | | | |

| User I/O Area | | | | | | | | | |
|-----------------------|----|------|--|----------------|----|------|--|---------------------|----|
| EMPNAME SMITH, JOE | | Ø↔→Ø | | EMPNO 12345 | | Ø↔→Ø | | ADDRESS NEW YORK | |
| 1 | 20 | | | 25 | 29 | | | 35 | 94 |
| Bytes | | | | | | | | | |

| Physical Segment on Disk After Update | | | | | | | | | |
|---------------------------------------|---|----------------|---|-----------------------|----|----------------|----|------|-----|
| LL 27 | | EMPNO 12345 | | EMPNAME SMITH, JOE | | BIRTHD Ø↔→Ø | | 0↔→0 | |
| 1 | 2 | 3 | 7 | 8 | 27 | 28 | 33 | 34 | 102 |
| Bytes | | | | | | | | | |

Figure 150. Second Example of a Missing Field on a Replace Call

The 'LL' field is maintained by IMS to include the full length of the ADDRESS field and all fields up to the ADDRESS field. BIRTHD is filled with blanks, because TYPE=C is specified on the FIELD statement in the DBD. Positions 34 to 42 are set to binary zeros, because the space was not defined by a FIELD statement in the DBD.

Inserting Missing Fields

When a variable-length segment is inserted into the database, the length field is set to the value of the minimum size needed to hold all sensitive fields.

- The 'LL' field is updated to include all sensitive fields.
- The TYPE= parameter on the DBD (see Figure 147 on page 232) determines the fill value for non-sensitive DBD fields.
- Binary zero is the fill value for space not defined by a FIELD statement in the DBD.

Figure 151 is an example of a missing field on an insert call using the DBD and PSB in Figure 147 on page 232.

Field-Level Sensitivity

| User I/O Area | | | | | |
|------------------------|-----|----------------|-----|--------------------|----|
| EMPNAME ADAMS, DICK | ␣↔␣ | EMPNO 23456 | ␣↔␣ | ADDRESS VERMONT | |
| 1 | 20 | 25 | 29 | 35 | 94 |
| Bytes | | | | | |

| Physical Segment on Disk After Insert | | | | | | | | | | | |
|---------------------------------------|----------------|------------------------|---------------|-----|--------------------|----|----|----|----|----|-----|
| LL 102 | EMPNO 23456 | EMPNAME ADAMS, DICK | BIRTHD ␣↔␣ | 0↔0 | ADDRESS VERMONT | | | | | | |
| 1 | 2 | 3 | 7 | 8 | 27 | 28 | 33 | 34 | 42 | 43 | 102 |
| Bytes | | | | | | | | | | | |

Figure 151. Example of a Missing Field on an Insert Call

The 'LL' field is maintained by IMS to include the full length of all sensitive fields up to and including the ADDRESS field. BIRTHD is filled with blanks, because TYPE=C was specified on the FIELD statement in the DBD. Positions 34 to 42 are set to binary zeros, because the space was not defined in a FIELD statement in the DBD.

Retrieving Partially Present Fields

If the last field in the physical variable-length segment at retrieval time is only partially present and if the data type is character (TYPE=C), data is returned to the application program padded with blanks on the right. Otherwise, the field is returned with a fill value based on the data type. Figure 152 is an example of a partially present field on a retrieval call based on the DBD and PSB in Figure 147 on page 232.

| Physical Segment on Disk | | | | | | | | | | | |
|--------------------------|----------------|-----------------------|---------------|---|---------------------|----|----|--|----|-----|--|
| LL 27 | EMPNO 12345 | EMPNAME SMITH, JOE | BIRTHD ␣↔␣ | | ADDRESS NEW YORK | | | | | | |
| 1 | 2 | 3 | 7 | 8 | 27 | 28 | 33 | | 43 | 102 | |
| Bytes | | | | | | | | | | | |

| User I/O Area | | | | | |
|-----------------------|-----|----------------|-----|---------------------|-----|
| EMPNAME SMITH, JOE | ␣↔␣ | EMPNO 12345 | ␣↔␣ | ADDRESS NEW YORK | ␣↔␣ |
| 1 | 20 | 25 | 29 | 35 | 94 |
| Bytes | | | | | |

Figure 152. Example of a Partially Present Field on a Retrieval Call

The ADDRESS field in the I/O area is padded with blanks to correspond to the length defined on the SEGM statement in the DBD.

Replacing Partially Present Fields

You should know the following information about replacing partially present fields:

- If segment length is increased on a REPL call, the field returned to the application program is written to the database if it has not been changed.
- If the data type of the field is character and the field is changed on a REPL call, the segment length is increased if necessary to include all non-blank characters in the changed data.
- If the data type is not character and the field is changed on a REPL call, the segment length is increased to contain the entire field.

Figure 153 is an example of a partially present field on a REPL call based on the DBD and PSB in Figure 147 on page 232.

| Physical Segment on Disk Before Update | | | | | | | | | |
|--|-------|------------|--------|---|----------|----|----|----|----|
| LL | EMPNO | EMPNAME | BIRTHD | | ADDRESS | | | | |
| 50 | 12345 | SMITH, JOE | 480207 | | NEW YORK | | | | |
| 1 | 2 | 3 | 7 | 8 | 27 | 28 | 33 | 43 | 50 |
| Bytes | | | | | | | | | |

| I/O Area | | | | | | | | | |
|------------|------|--|-------|------|--|----------|--|----|--|
| EMPNAME | ↔↔↔↔ | | EMPNO | ↔↔↔↔ | | ADDRESS | | | |
| SMITH, JOE | | | 12345 | | | NEW YORK | | | |
| 1 | 20 | | 25 | 29 | | 35 | | 94 | |
| Bytes | | | | | | | | | |

| Physical Segment on Disk After Update | | | | | | | | | |
|---------------------------------------|-------|------------|--------|---|----------|----|----|----|----|
| LL | EMPNO | EMPNAME | BIRTHD | | ADDRESS | | | | |
| 52 | 12345 | SMITH, JOE | 480207 | | NEW YORK | | | | |
| 1 | 2 | 3 | 7 | 8 | 27 | 28 | 33 | 43 | 52 |
| Bytes | | | | | | | | | |

Figure 153. Example of a Partially Present Field on a REPL Call

The 'LL' field is changed from 50 to 52 by DL/I to accommodate the change in the field length of ADDRESS.

General Considerations for Using Field-Level Sensitivity

- Field-level sensitivity is not supported for GSAM, MSDB, or DEDB databases.
- Fields referenced in PSBGEN with SENFLD statements must be defined in DBDGEN with FIELD statements.
- The same DBD field can be referenced in more than one SENFLD statement.
- When using field-level sensitivity, the application program always sees a fixed length segment for a given PCB, regardless of whether the segment is fixed or variable.
- Application programs must be sensitive to any field referenced in an SSA, except the sequence field.
- Application programs must be sensitive to the sequence field, if present, for insert or load.
- Field-level sensitivity and segment level sensitivity can be mixed in the same PCB.
- Non-referenced, non-defined fields are set to binary zeros as fill characters, when required, during insert or replace operations.
- Using call/trace with the compare option increases the amount of storage required in the PSB work pool.

Multiple Data Set Groups

The following database types support multiple data set groups:

- SHISAM
- HDAM
- PHDAM
- HIDAM

Multiple Data Set Groups

- PHIDAM

Although this book has explored storing a database on a single or a single pair of data sets, HD databases can be stored on more than the one or two data sets required for database storage. You have seen that an HD database is stored on an ESDS, if VSAM is being used, or an OSAM data set, if OSAM is being used.

HD databases can be stored on multiple data sets. When storing a database on multiple data sets, the terms primary and secondary data set group are used to distinguish between the one or more data sets that must be specified for the database (called the primary data set group) and the one or more data sets you are allowed to specify for the database (called secondary data set groups).

In HD databases, a single data set is used for storage rather than a pair of data sets. The primary data set group therefore consists of the ESDS (if VSAM is being used) or OSAM data set (if OSAM is being used) on which you must specify storage for your database. The secondary data set group is an additional ESDS or OSAM data set on which you are allowed to store your database.

As many as ten data set groups can be used in HISAM and HD databases, that is, one primary data set group and a maximum of nine secondary data set groups.

Why Use Multiple Data Set Groups?

When you design database records, you design them to meet the processing requirements of many applications. You decide what segments will be in a database record and their hierarchic sequence within a database record. These decisions are based on what works best for all of your application program's requirements. However, the way in which you arranged segments in a database record no doubt suits the processing requirements of some applications better than others. For example, look at the two database records shown in Figure 154. Both of them contain the same segments, but the hierarchic sequence of segments is different.

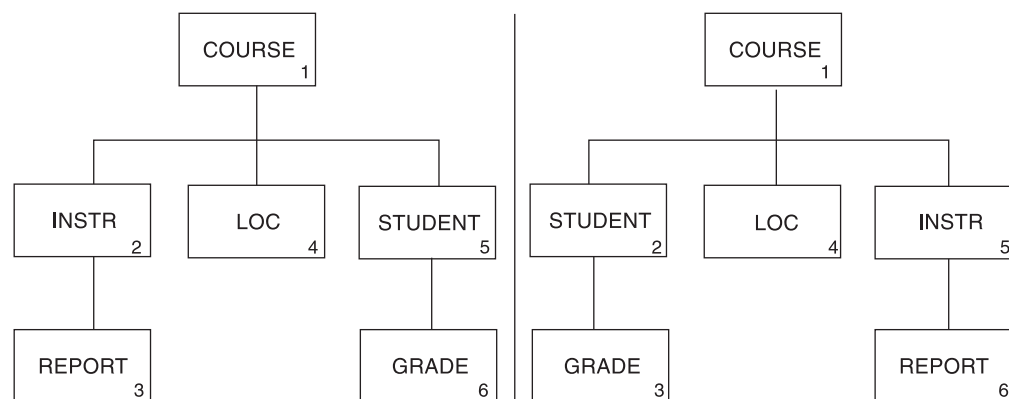


Figure 154. Hierarchy of Applications That Need to Access INSTR and LOC Segments

The hierarchy on the left favors applications that need to access INSTR and LOC segments. The hierarchy on the right favors applications that need to access STUDENT and GRADE segments. (Favor, in this context, means that access to the segments is faster.) If the applications that access the INSTR and LOC segments are more important than the ones that access the STUDENT and GRADE segments, you can use the database record on the left. But if both applications are equally important, you can split the database record into different data set groups. This will give both types of applications good access to the segments each needs.

To split the database record, you would use two data set groups. As shown in Figure 155, the first data set group contains the COURSE, INSTR, REPORT, and LOC segments. The second data set group contains the STUDENT and GRADE segments.

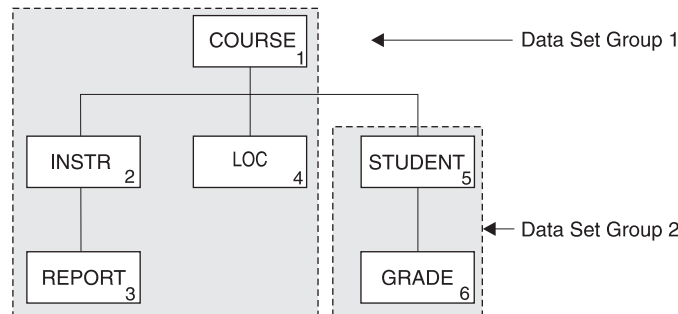


Figure 155. Database Record Split into Two Database Groups

Other uses of multiple data set groups include:

- Separating infrequently-used segments from high-use segments.
- Separating segments that frequently have information added to them from those that do not. For the former segments, you might specify additional free space so conditions are optimum for additions.
- Separating segments that are added or deleted frequently from those that are not. This can keep space from being fragmented in the main database.
- Separating segments whose size varies greatly from the average segment size. This can improve use of space in the database. Remember, the bit map in an HD database indicates whether space is available for the *longest* segment type defined in the data set group. It does not keep track of smaller amounts of space. If you have one or more segment types that are large, available space for smaller segments will not be utilized, because the bit map does not track it.

HD Databases Using Multiple Data Set Groups

The following rules must be followed when using a multiple data set group in an HD database:

- As many as ten data set groups can be defined.
- The root segment in a database record must be in the primary data set group.

In the database record shown in Figure 156 on page 238, segments COURSE (1), INSTR (2), LOC (4), and STUDENT (5) could go in one data set group, while segments REPORT (3) and GRADE (6) could go in a second data set group. Examples of how this HD database record could be divided into three groups are in Table 18.

Table 18. Examples of Multiple Data Set Grouping

| Data Set Group 1 | Data Set Group 2 | Data Set Group 3 |
|----------------------|----------------------|------------------|
| Segment 1 | Segments 2, 5, and 6 | Segments 3 and 4 |
| Segments 1, 3, and 6 | Segments 2 and 5 | Segment 3 |
| Segments 1, 3, and 6 | Segments 2 and 5 | Segment 4 |

Multiple Data Set Groups

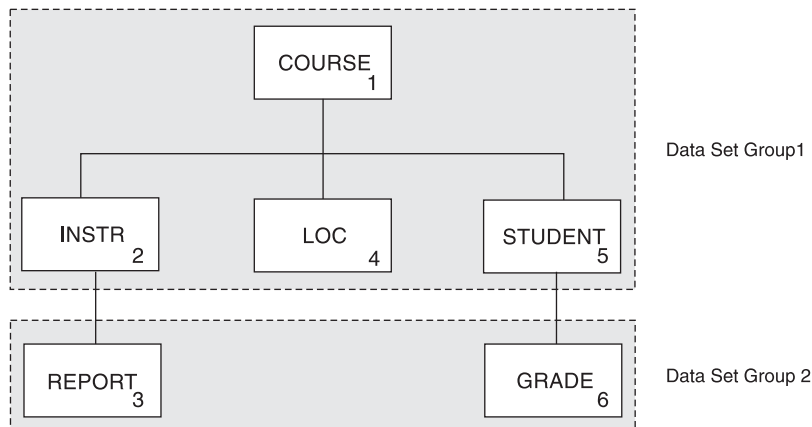


Figure 156. Example of How to Divide an HD Database Record

- Segments separated into different data set groups must be connected by physical child first pointers. For example, in Figure 157 the INSTR segment in the primary data set group must point to the first occurrence of its physical child REPORT in the secondary data set group, and STUDENT must point to GRADE.

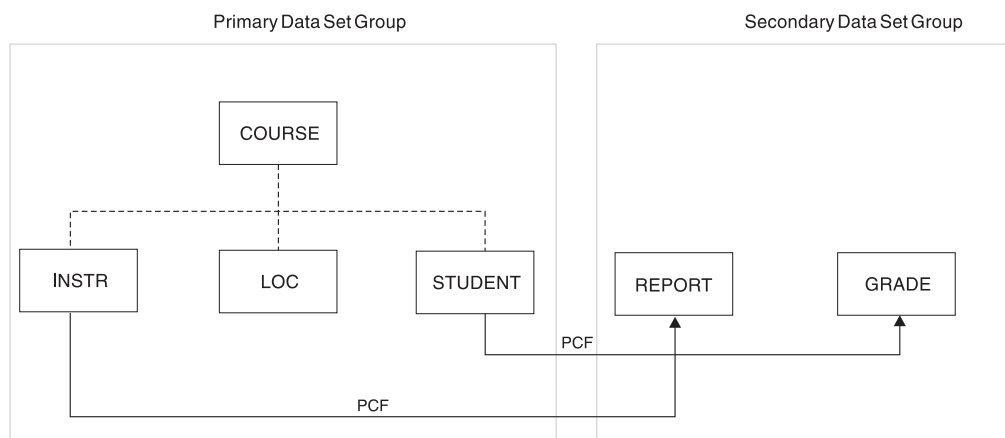


Figure 157. Connecting Segments in Multiple Data Set Groups Using Physical Child First Pointers

How HD Records Are Stored in Multiple Data Set Groups

Now that you have seen what segments can be stored in a single data set group in an HD database, this topic looks at how segments are stored. Figure 158 on page 239 shows one database record:

- Stored in an HDAM or a PHDAM database using two data set groups
- Stored in a HIDAM or a PHIDAM database using two data set groups

Specify in the DBD which segment types need to be put in a data set group. Based on that information, IMS automatically loads segments into the correct data set group. In this example, the user specified that four segment types in the database record were put in the primary data set group (COURSE, INSTR, LOC, STUDENT) and two segment types were put in the secondary data set group (REPORT, GRADE).

In the HDAM or PHDAM database, note that only the primary data set group has a root addressable area. The secondary data set group is additional overflow storage.

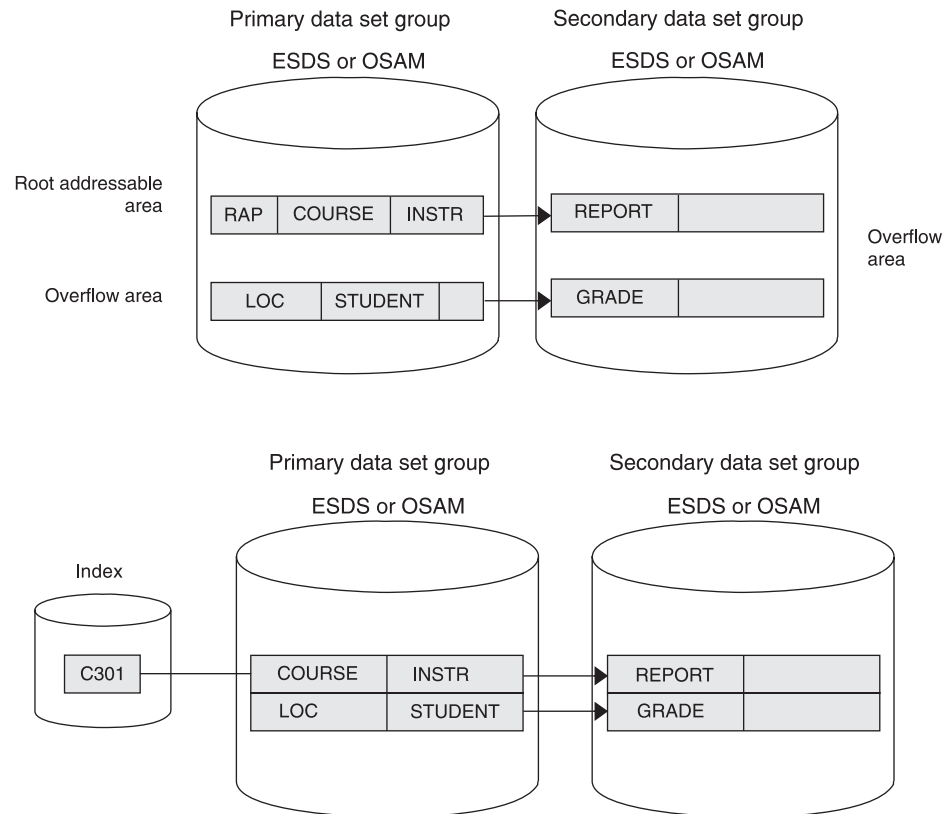


Figure 158. HD Database Record in Storage When Multiple Data Set Groups Are Used

Specifying Use of Multiple Data Set Groups in HD and PHD Databases

You can specify multiple data set groups to IMS in the DBD. For HDAM databases, use the DATASET statement. For PHDAM databases, use the DSGROUP parameter in the SEGM statement. You can group the segments any way, but you still must list the segments in hierarchical sequence in the DBD.

The following examples use the database record used in “Why Use Multiple Data Set Groups?” on page 236 and “HD Databases Using Multiple Data Set Groups” on page 237. The first example, Figure 159, shows two groups: data set group A contains COURSE and INSTR, data set group B contains all of the other segments. The second example shows a different grouping. Note the differences in DBDs when the groups are not in sequential hierarchical order of the segments.

Multiple Data Set Groups

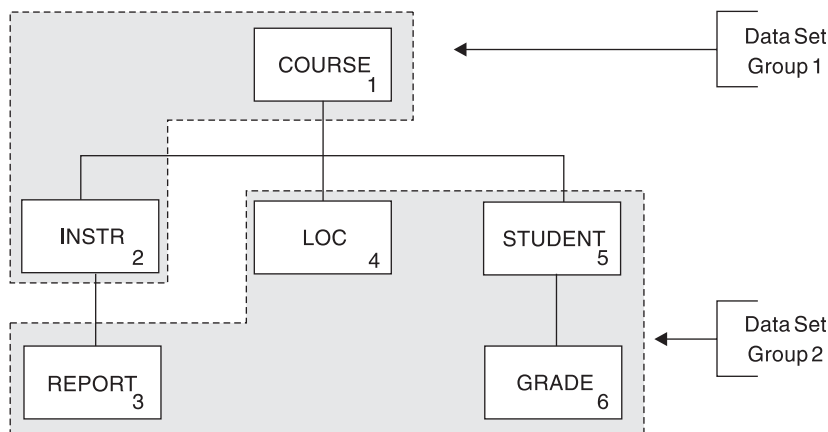


Figure 159. First Example of Data Set Groups

Figure 160 is the HDAM DBD for the first example. Note that the segments are grouped by the DATASET statements preceding the SEGM statements and that the segments are listed in hierarchical order. In each DATASET statement, the DD1= parameter names the VSAM ESDS or OSAM data set that will be used. Also, each data set group can have its own characteristics, such as device type.

```

DBD      NAME=HDMDSG,ACCESS=HDAM,RMNAME=(DFSHDC40,8,500)
DSA      DATASET DD1=DS1DD,
SEGM     NAME=COURSE,BYTES=50,PTR=T
FIELD    NAME=(CODCOURSE,SEQ),BYTES=10,START=1
SEGM     NAME=INSTR,BYTES=50,PTR=T,PARENT=((COURSE,SNGL))
DSB      DATASET DD1=DS2DD,DEVICE=2314
SEGM     NAME=REPORT,BYTES=50,PTR=T,PARENT=((INSTR,SNGL))
SEGM     NAME=LOC,BYTES=50,PTR=T,PARENT=((COURSE,SNGL))
SEGM     NAME=STUDENT,BYTES=50,PTR=T,PARENT=((COURSE,SNGL))
SEGM     NAME=GRADE,BYTES=50,PTR=T,PARENT=((STUDENT,SNGL))
DBDGEN

```

Figure 160. HDAM DBD for First Example of Data Set Groups

Figure 161 shows the DBD for a PHDAM database. Instead of using the DATASET statement, use the DSGROUP parameter in the SEGM statement. The first two segments do not have DSGROUP parameters because it is assumed that they are in the first group.

```

DBD      NAME=HDMDSG,ACCESS=PHDAM,RMNAME=(DFSHDC40,8,500)
SEGM     NAME=COURSE,BYTES=50,PTR=T
FIELD    NAME=(CODCOURSE,SEQ),BYTES=10,START=1
SEGM     NAME=INSTR,BYTES=50,PTR=T,PARENT=((COURSE,SNGL))
SEGM     NAME=REPORT,BYTES=50,PTR=T,PARENT=((INSTR,SNGL)),DSGROUP=B
SEGM     NAME=LOC,BYTES=50,PTR=T,PARENT=((COURSE,SNGL)),DSGROUP=B
SEGM     NAME=STUDENT,BYTES=50,PTR=T,PARENT=((COURSE,SNGL)),DSGROUP=B
SEGM     NAME=GRADE,BYTES=50,PTR=T,PARENT=((STUDENT,SNGL)),DSGROUP=B
DBDGEN

```

Figure 161. PHDAM DBD for First Example of Data Set Groups

The second example, Figure 162 on page 241, differs from the first example in that the groups do not follow the order of the hierarchical sequence. The segments must be listed in the DBD in hierarchical sequence, so additional DATASET statements or DSGROUP parameters are required.

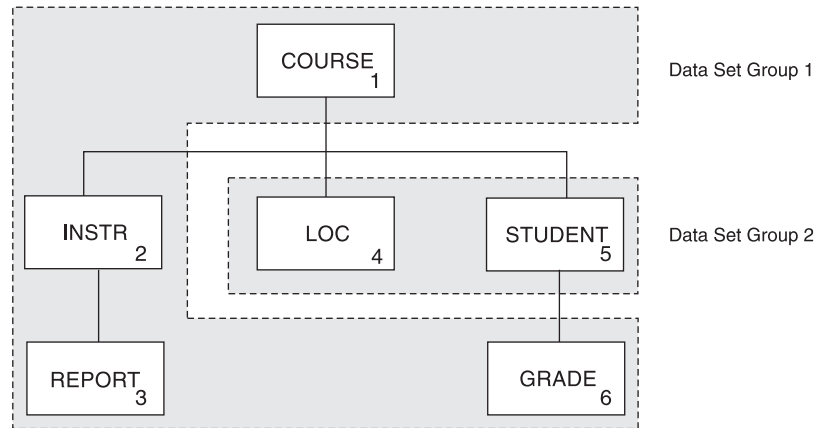


Figure 162. Second Example of Data Set Groups

Figure 163 is the DBD for an HDAM database of the second example. It is similar to the first example, except that because the sixth segment is part of the first group, you need another DATASET statement before it with the DSA label. The additional DATASET label groups the sixth segment with the first three.

```

DBD    NAME=HDMDSG,ACCESS=HDAM,RMNAME=(DFSHDC40,8,500)
DSA    DATASET DD1=DS1DD,
SEGM   NAME=COURSE,BYTES=50,PTR=T
FIELD  NAME=(CODCOURSE,SEQ),BYTES=10,START=1
SEGM   NAME=INSTR,BYTES=50,PTR=T,PARENT=((COURSE,SNGL))
SEGM   NAME=REPORT,BYTES=50,PTR=T,PARENT=((INSTR,SNGL))
DSB    DATASET DD1=DS2DD,DEVICE=2314
SEGM   NAME=LOC,BYTES=50,PTR=T,PARENT=((COURSE,SNGL))
SEGM   NAME=STUDENT,BYTES=50,PTR=T,PARENT=((COURSE,SNGL))
DSA    DATASET DD1=DS1DD
SEGM   NAME=GRADE,BYTES=50,PTR=T,PARENT=((STUDENT,SNGL))
DBDGEN
  
```

Figure 163. HDAM DBD for Second Example of Data Set Groups

Figure 164 is the DBD for a PHDAM database of the second example. It is similar to the first example, except that because the sixth segment is part of the first group, you must explicitly group it with the first two segments by using the DSGROUP parameter.

```

DBD    NAME=HDMDSG,ACCESS=PHDAM,RMNAME=(DFSHDC40,8,500)
SEGM   NAME=COURSE,BYTES=50,PTR=T
FIELD  NAME=(CODCOURSE,SEQ),BYTES=10,START=1
SEGM   NAME=INSTR,BYTES=50,PTR=T,PARENT=((COURSE,SNGL))
SEGM   NAME=REPORT,BYTES=50,PTR=T,PARENT=((INSTR,SNGL)),
SEGM   NAME=LOC,BYTES=50,PTR=T,PARENT=((COURSE,SNGL)),DSGROUP=B
SEGM   NAME=STUDENT,BYTES=50,PTR=T,PARENT=((COURSE,SNGL)),DSGROUP=B
SEGM   NAME=GRADE,BYTES=50,PTR=T,PARENT=((STUDENT,SNGL)),DSGROUP=A
DBDGEN
  
```

Figure 164. PHDAM DBD for Second Example of Data Set Groups

VSAM KSDS CI Reclaim For Full-Function Databases

In data-sharing or XRF environments, IMS can reclaim the storage used for empty VSAM KSDS control intervals (CIs) in full-function databases by using a process called CI reclaim. CI reclaim enhances the performance of database GN, GU, and ISRT calls by reducing the number of empty CIs that VSAM reads on subsequent DL/I calls.

CI Reclaim

During the deletion of the last record in a CI, IMS flags the CI for reclamation. After the deletion has been committed, IMS reclaims the CI the next time the CI is read for a DL/I call that has update access. On a single DL/I call, IMS reclaims every empty CI flagged for reclamation that it reads until the DL/I call has been satisfied.

CI reclaim is not a replacement for the routine reorganization of KSDS data sets.

CI reclaim works only under the following circumstances:

- You specify ERASE=YES, or accepted it as the default, on the DBD statement of the DFSVSMxx PROCLIB member.
- Your KSDS uses unique keys.
- The application program issuing the GN, GU, or ISRT call has update access to the database. Using a processing option PROCOPT of "I", "R", "D" or "A" guarantees that the dataset is open for update.

Restriction: CI reclaim does not occur for SHISAM databases. When a large number of records in a SHISAM database are deleted, particularly a large number of consecutive records, serious performance degradation can occur. Eliminate empty CIs and resolve the problem by using VSAM REPRO.

Related Reading:

- For information about the data sharing and XRF environments, see *IMS Version 9: Administration Guide: System*.
- For information about the DFSVSMxx PROCLIB member, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.
- For information about the VSAM REPRO command, see *DFSMS Access Method Services for Catalogs*.

HALDB Single Partition Processing

BMP, JBP, and batch-processing applications can process a single partition of a HALDB database independently of the rest of the HALDB. The partition independence is similar to the independent processing of partitions by utilities. To restrict processing to a single partition, restrict DB PCB usage by specifying the label name of the DB PCB or the nth position of the DB PCB, and the partition name in the HALDB control statements.

Related Reading: For information on specifying single partition processing, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Logical Relationships in Single Partition Processing

An application can process single partitions with logical relationships. If a logical child is in the single partition that the application has access to, and its logical parent is in another partition, the application can process the logical parent, even though it is in another partition. Because of a logical relationship, an application with restricted access can process a partition that it does not have direct access to.

Secondary Indexes in Single Partition Processing

You can restrict application programs to a single partition of a HALDB partitioned secondary index (PSINDEX) by specifying the name of the PSINDEX partition in the DFSHALDB statement.

You can use single partition processing with a PSINDEX regardless of whether your application program processes your PSINDEX as a standalone database or as an alternate processing sequence for a PHDAM or PHIDAM database.

The partitions in a PSINDEX do not correspond to the partitions in the PHDAM or PHIDAM database that the PSINDEX indexes. Consequently, when you specify single partition processing for a PSINDEX, single partition processing only applies to the PSINDEX partition, not the partitions of the PHDAM or PHIDAM database. The target segments can be in any partition in the indexed PHDAM or PHIDAM database.

Similarly, if you specify single partition processing for a PHDAM or PHIDAM database, the single partition processing does not apply to the partitions in any associated PSINDEXs.

Regardless of whether you are using single partition processing with a PSINDEX or with an indexed PHDAM or PHIDAM database, single partition processing does not affect the internal updating of a secondary index by IMS when a target segment is updated. For example, if an application program is restricted to a single partition of the PSINDEX and inserts a segment into the indexed PHDAM or PHIDAM database, the corresponding new index entry can be inserted in any partition of the PSINDEX.

Partition Selection

A partition is selected by using the root key for the DL/I call and the high key defined for the partition. When access is restricted to a single partition and the root key is outside the key range of the partition, status code FM or GE is returned.

If you use a partition selection exit routine, the routine is called when the DL/I call provides a specific root key. The exit routine selects a partition based on the root key given. If the partition selected is different than the one that the application has access to, status code FM or GE is returned. The exit routine is not called to select a first partition or next partition.

When access is restricted to a single partition, the first partition is always the partition to which access is restricted, and the next partition does not exist.

Recommendation: If restricting processing to a single partition, the SSA should include only the root keys that are in the key range of the partition.

Examples of Single Partition Processing

For the following examples, the DB PCB usage is restricted to HALDB partition 2, which contains the records with root keys 201 through 400.

GU rootkey=110

The root key 110 is outside the range of root keys for the partition. FM status code is returned.

GU rootkey=240 GN rootkey=110

Moves forward from root key 240 to find key equal to 110. Because 110 is lower than 240, GE status code is returned.

GU rootkey=240 GN rootkey>=110

Moves forward from root key 240 to find key equal to or greater than 110. If key not found before reaching end of partition, GB status code is returned.

GN rootkey>=110

Attempts to start search at key 110. Because key is outside root key range of partition, FM status code is returned.

Integrated HALDB Online Reorganization Function

With the integrated HALDB Online Reorganization (OLR) function, you can reorganize HALDB partitions online, improving database performance without disrupting data availability. HALDB OLR can reorganize any number of partitions, singly or in parallel.

Related Reading: For complete information on the online reorganization of HALDBs, see “HALDB Online Reorganization” on page 370.

Storing XML Data in IMS Databases

You can store and retrieve XML documents in IMS databases using Java application programs. When storing and retrieving XML documents, the XML documents must be valid to XML schemas generated by the DLIModel utility. The XML schemas must match the hierarchical structure of the IMS database.

XML documents can be stored in IMS databases using any combination of two storage methods to best fit the structure of the XML document:

Decomposed XML storage

The XML tags are removed from the XML document and only the data is extracted. The extracted data is converted into traditional IMS field types and inserted into the database. Use this approach in the following scenarios:

- XML applications and non-XML applications must access the same database.
- Extensive searching of the database is needed.
- A strict XML schema is available.

Intact XML storage

The XML document is stored, with its XML structure and tags intact, in an IMS database designed exclusively for storing intact XML documents. In this case, only Java application programs can access the data in the database. Because the XML document does not have to be regenerated when the data is retrieved from the database, the retrieval of the XML data is typically faster than when it is stored without its XML tagging. Use this approach in the following scenarios:

- Faster storage and retrieval of XML documents are needed.
- Less searching of the database is required.
- The XML schema requires more flexibility.

Related Reading:

- For more information about the DLIModel utility, see *IMS Version 9: Utilities Reference: System*.
- For more information about storing XML data in IMS databases, see *IMS Version 9: IMS Java Guide and Reference*.

Chapter 9. Designing Full-Function Databases

After you determine the type of database and optional functions that best suit your application's processing requirements, you need to make a series of decisions about database design and use of options. This set of decisions primarily determines how well your database performs and how well it uses available space. This series of decisions is made based on:

- The type of database and optional functions you have already chosen
- The performance requirements of your applications
- How much storage you have available for use online

In this Chapter:

- "Specifying Free Space (HDAM, PHDAM, HIDAM, and PHIDAM Only)"
- "Estimating the Size of the Root Addressable Area (HDAM or PHDAM Only)" on page 246
- "Determining Which Randomizing Module to Use (HDAM and PHDAM Only)" on page 247
- "Choosing HDAM or PHDAM Options" on page 248
- "Choosing a Logical Record Length for a HISAM Database" on page 249
- "Choosing a Logical Record Length for HD Databases" on page 252
- "Determining the Size of CIs and Blocks" on page 252
- "Buffering Options" on page 253
- "OSAM Sequential Buffering" on page 257
- "VSAM Options" on page 264
- "OSAM Options" on page 269
- "Dump Option (DUMP Parameter)" on page 269
- "Deciding Which FIELD Statements to Code in the DBD" on page 269
- "Planning for Maintenance" on page 269

Specifying Free Space (HDAM, PHDAM, HIDAM, and PHIDAM Only)

As you have seen, dependent segments inserted after an HD database is loaded are put as close as possible to the segments to which they are related. (When segments are close to the segments that point to them, the I/O time needed to retrieve a dependent segment is shorter. The I/O time is shorter because the seek time and rotational delay time are shorter.) However, as the database grows and available space decreases, dependent segments are increasingly put further from their related segments. When this happens, performance decreases, a problem that can only be eliminated by reorganizing the database.

To minimize the effect of insert operations after the database is loaded, allocate free space in the database when it is initially loaded. Free space allocation in the database will reduce the performance impact caused by insert operations, and therefore, decrease the frequency with which HD databases must be reorganized.

For OSAM data sets and VSAM ESDS, free space is specified in the FRSPC= keyword of the DATASET statement in the DBD. In the keyword, one or both of the following operands can be specified:

- Free block frequency factor (fbff). The fbff specifies that every nth block or CI in a data set group be left as free space when the database is loaded (where fbff=n).

Specifying Free Space (HDAM, PHDAM, HIDAM, and PHIDAM Only)

The range of fbff includes all integer values from 0 to 100, except 1. Avoid specifying fbff for HDAM or PHDAM databases. If you specify fbff for HDAM or PHDAM databases and if at load time the randomizing module generates the relative block or CI number of a block or CI marked as free space, the randomizer must store the root segment in another block.

If you specify fbff, every nth block or CI will be considered a second-most desirable block or CI by the HD Space Search Algorithm. This is true unless you specify SEARCHA=1 in the DATASET macro of the DBDGEN utility. By specifying SEARCHA=1, you are telling IMS not to search for space in the second-most desirable block or CI.

Related Reading:

- For details on the HD Space Search Algorithm, see “How the HD Space Search Algorithm Works” on page 107.
- For more information on the SEARCHA keyword, see *IMS Version 9: Utilities Reference: Database and Transaction Manager*.
- Free space percentage factor (fspf). The fspf specifies the minimum percentage of each block or CI in a data set group to be left as free space when the database is loaded. The range of fspf is from 0 to 99.

Note: This free space applies to VSAM ESDS and OSAM data sets. It does not apply to HIDAM or PHIDAM index databases or to DEDBs.

For VSAM KSDS, free space is specified in the FREESPACE parameter of the DEFINE CLUSTER command. This VSAM parameter is disregarded for a VSAM ESDS data set used for HIDAM, PHIDAM, HDAM, or PHDAM. This command is explained in detail in *DFSMS Access Method Services for Catalogs*.

Estimating the Size of the Root Addressable Area (HDAM or PHDAM Only)

To estimate the size of the root addressable area, use the following formula:

$$(A \times B) / C = D$$

where:

- A =** the number of bytes of a database record to be stored in the root addressable area
- B =** the expected number of database records
- C =** the number of bytes available for data in each CI or block CI or block size, minus overhead)
- D =** the size you will need, in blocks or CIs, for the root addressable area.

If you have specified free space for the database, include it in your calculations for determining the size of the root addressable area. Use the following formula to accomplish this step:

$$(D \times E \times G) / F = H$$

where:

- D =** the size you calculated in the first formula (the necessary size of the root addressable area in block or CIs)
- E =** how often you are leaving a block or CI in the database empty for free space (what you specified in the fbff operand in the DBD)

F = (E-1) (fbff-1)

G = 100 - 100 - fspf The fspf is the minimum percentage of each block or CI you are leaving as free space (what you specified in the fspf operand in the DBD)

H = the total size you will need, in blocks or CIs

Specify the number of blocks or CIs you need in the root addressable area in the RMNAME=rbn keyword in the DBD statement in the DBD.

Determining Which Randomizing Module to Use (HDAM and PHDAM Only)

As you have seen, a randomizing module is required to store and access HDAM or PHDAM database records. This module converts the key of a root segment to a relative block number and RAP number. These numbers are then used to store or access HDAM or PHDAM root segments. An HDAM database or a PHDAM partition uses only one randomizing module, but several databases and partitions can share the same module. Four randomizing modules are supplied with IMS.

Normally, one of the four randomizing modules supplied with the system will work for your database. These modules, and the arithmetic techniques they use, are described in detail in *IMS Version 9: Customization Guide*.

Partition selection is completed prior to invoking the randomizing module on PHDAM databases. The randomizing module selects locations only within a partition.

Write Your Own Randomizing Module

If, given your root key distribution, none of these randomizing modules works well for you, write your own randomizing module. If you write your own randomizing module, one of your goals is to have it distribute root segments so that, when subsequently accessing them, only one read and one seek operation is required. When a root key is given to the randomizing module, if the relative block number the randomizer produces is the block actually containing the root, only one read and seek operation is required (access is fast). The randomizing module you write should allow you to vary the number of blocks and RAPs you specify, so blocks and RAPs can be used for tuning the system. The randomizing module should also distribute roots randomly, not randomize to bit map locations, and keep packing density high. *IMS Version 9: Customization Guide* tells you what the interface to your randomizing module should be.

Assess the Effectiveness of the Randomizing Module

One way to determine the effectiveness of a given randomizing module for your database is to run the IMS High Performance Pointer Checker (HD Tuning Aid). This tool produces a report in the form of a map showing how root segments are stored in the database. It shows you root segment storage based on the number of blocks or CIs you specified for the root addressable area and the number of RAPs you specified for each block or CI. By running the HD Tuning Aid against the various randomizing modules, you can see which module gives you the best distribution of root keys in your database. In addition, by changing the number of RAPs and blocks or CIs you specify, you can see (given a specific randomizing module) which combination of RAPs and blocks or CIs produces the best root segment distribution.

Determining Which Randomizing Module To Use

Before choosing a randomizing module, read “Adjusting HDAM and PHDAM Options” on page 411, which discusses how you can adjust HDAM or PHDAM options, including the randomizing module, to tune your database after it is running.

Choosing HDAM or PHDAM Options

In an HDAM or a PHDAM database, the options you choose can greatly affect performance. The options discussed here are those you specify in the RMNAME keyword in the DBD statement or when using the HALDB Partition Definition utility. Figure 165 shows the format for specifying the RMNAME parameter. The definition list that follows explains the meaning of *mod*, *anch*, *rbn*, and *bytes*.

RMNAME=(*mod,anch,rbn,bytes*)

Figure 165. Specifying the RMNAME keyword

| | |
|--------------|--|
| <i>mod</i> | Name of the randomizing module you have chosen |
| <i>anch</i> | Number of RAPs in a block or CI |
| <i>rbn</i> | Number of blocks or CIs in the root addressable area |
| <i>bytes</i> | Maximum number of bytes of a database record to be put in the root addressable area when segments in the database records are inserted consecutively (without intervening processing operations) |

Minimizing I/O Operations

In choosing these HDAM or PHDAM options, your primary goal is to minimize the number of I/O operations it takes to access a database record or segment. The fewer I/O operations, the faster the access time. Performance is best when:

- The number of RAPs in a block or CI is equal to the number of roots in the block or CI (block or CI space is not wasted on unused RAPs).
- Unique block and RAP numbers are generated for most root segments (thereby eliminating long synonym chains).
- Root segments are stored in key sequence.
- All frequently used dependent segments are in the root addressable area (access to the root addressable area is faster than access to the overflow area) and in the same block or CI as the root.

Your choice of a randomizing module (discussed in “Determining Which Randomizing Module to Use (HDAM and PHDAM Only)” on page 247) determines how many addresses are unique for each root and whether roots are stored in key sequence. In general, a randomizing module is considered efficient if roots are distributed evenly in the root addressable area. You can experiment with different randomizing modules. Try various combinations of the *anch*, *rbn*, and *bytes* operands to see what effect they have on distribution of root segments.

Maximizing Packing Density

A secondary goal in choosing HDAM or PHDAM options is to maximize packing density without adversely affecting performance. Packing density is the percentage of space in the root addressable area being used for root segments and the dependent segments associated with them. Packing density is determined as follows:

Packing density =
(Number of roots x root bytes) /
(Number of CIs in the root addressable area x Usable space in the CI)

root bytes

The average number of bytes in each root in the root addressable area.

Usable space in the CI

The CI or block size minus (as applicable) space for the FSEAP, RAPs, VSAM CIDE, VSAM RDE, and free space.

Packing density should be high, but, as the percentage of packing density increases, the number of dependent segments put into overflow storage can increase. In addition, performance for processing of dependent segments decreases when they are in overflow storage. All of the operands you can specify in the RMNAME= keyword affect packing density. So, to optimize packing density, try different randomizing modules and various combinations of the anch, rbn, and bytes operands.

Choosing a Logical Record Length for a HISAM Database

In a HISAM database, your choice of a logical record length is important because it can affect both the access time and the use of space in the database. The relative importance of each depends on your individual situation. To get the best possible performance and an optimum balance between access time and the use of space, plot several trial logical record lengths and test them before making a final choice.

Logical Record Length Considerations

The following should be considered:

- Only complete segments can be stored in a logical record. Therefore, the space between the last segment that fit in the logical record and the end of the logical record is unused.
- Each database record starts at the beginning of a logical record. The space between the end of the database record and the end of the last logical record containing it is unused. This unused space is relative to the average size of your database records.
- Very short or very long logical records tend to increase wasted space. If logical records are short, the number of areas of unused space increases. If logical records are long, the size of areas of unused space increases. Figure 166 shows why short or long logical records increase wasted space.

Choose a logical record length that minimizes the amount of unused space at the end of logical records.

The database record shown in Figure 166 on page 250 is stored on three short logical records in Figure 167 on page 250 and in two longer logical records in Figure 168 on page 250. Note the three areas of unused space.

Choosing a Logical Record Length for a HISAM Database




Figure 166. Database Record for Logical Record Examples

In Figure 167, note the three areas of unused space. In Figure 168, there are only two areas of unused space, rather than three, but the total size of the areas is larger.




Figure 167. Short Logical Records




Figure 168. Long Logical Records

Segments in a database record that do not fit in the logical record in the primary data set are put in one or more logical records in the overflow data set. More read and seek operations, and therefore longer access time, are required to access logical records in the overflow data set than in the primary data set. This is especially true as the database grows in size and chains of overflow records develop. Therefore, you should try to put the most-used segments in your database record in the primary data set. When choosing a logical record length the primary data set should be as close to average database record length as possible. This results in a minimum of overflow logical records and thereby minimizes performance problems. When you calculate the average record length, beware of unusually long or short records that can skew the results.

A read operation reads one CI into the buffer pool. CIs contain one or more logical records in a database record. Because of this, it takes as many read and seek operations to access an entire database record as it takes CIs to contain it. In Figure 170 on page 251, each CI contains two logical records, and two CIs are required to contain the database record shown in Figure 169 on page 251. Consequently, it takes two read operations to get these four logical records into the buffer.

Choosing a Logical Record Length for a HISAM Database




Figure 169. Database Record for Logical Records Example




Figure 170. Logical Records Example with Two Read Operations

The number of read and seek operations required to access a database record increases as the size of the logical record decreases. The question to consider is: Do you often need access to the entire database record? If so, you should try to choose a logical record size that will usually contain an entire database record. If, however, you typically access only one or a few segments in a database record, choice of a logical record size large enough to contain the average database record is not as important.

Consider what will happen in the following setup example in which you need to read database records, one after another:

- Your CI or block size is 2048 bytes.
- Your Logical record size is 512 bytes.
- Your Average database record size is 500 bytes.
- The range of your database record sizes is 300 to 700 bytes.

Because your logical and average database record sizes are about equal (512 and 500), approximately one of every two database records will be read into the buffer pool with one read operation. (This assumption is based on the average size of database records.) If, however, your logical record size were 650, you would access most database records with a single read operation. An obvious trade-off exists here, one you must consider in picking a logical record length for HISAM data sets. If your logical record size were 650, much unused space would exist between the end of an average database record and the last logical record containing it.

Rules to Observe

The following rules must be observed when choosing a logical record length for HISAM data sets:

- Logical record size in the primary data set must be at least equal to the size of the root segment, plus its prefix, plus overhead. If variable-length segments are

Choosing a Logical Record Length for a HISAM Database

used, logical record size must be at least equal to the size of the longest root segment, plus its prefix, plus overhead. Five bytes of overhead is required for VSAM.

- Logical record size in the overflow data set must be at least equal to the size of the longest segment in the overflow data set, plus its prefix, plus overhead. Five bytes of overhead is required for VSAM.
- Logical record lengths in the overflow data set must be equal to or greater than logical record length in the primary data set.
- The maximum logical record size is 30720 bytes.
- Except for SHISAM databases, logical record lengths must be an even number.

Calculating How Many Logical Records Are Needed to Hold a Database Record

Calculate the average size of a database record before plotting various logical record sizes. By calculating the average size of a database record, given a specific logical record size, you can see how many logical records it takes to hold a database record (of average size).

Related Reading: To determine the average size of your database records, see “Estimating the Minimum Size of the Database” on page 315.

Specifying Logical Record Length

Specify the length of the logical records in the RECORD= operand of the DATASET statement in the DBD.

Choosing a Logical Record Length for HD Databases

In HD databases, the important choice is not logical record length but CI or block size. Logical record length is the same as block size when VSAM is used. Logical record size is equal to CI size, minus 7 bytes of overhead (4 bytes for a CIDE, 3 bytes for an RDF).

Related Reading: See “Determining the Size of CIs and Blocks” for information on determining CI or block size.

As with HISAM databases, specify the length of the logical records in the RECORD= operand of the DATASET statement in the DBD.

Determining the Size of CIs and Blocks

You can specify the DEDB CI resource size for your database. (If you do not specify it, the DBDGEN utility will calculate it for you.) Based on CI size, VSAM determines the size of physical blocks on a DASD track. VSAM always uses the largest possible physical block size, because the largest block size best utilizes space on the track. So your choice of a CI size is an important one. Your goal in picking it is to keep a high percentage of space on the track for your data, rather than for device overhead.

Track sizes vary from one device to another, and many different CI sizes you can specify exist. Because you can specify different CI sizes, the physical block size that VSAM picks varies and is based on device overhead factors. For information about using VSAM data sets, refer to *DFSMS Access Method Services for Catalogs*.

Buffering Options

Database buffers are defined areas in virtual storage. When an application program processes a segment in the database, the entire block or CI containing the segment is read from the database into a buffer. The application program processes the segment while it is in the buffer. If the processing involves modifying any segments in the buffer, the contents of the buffer must eventually be written back to the database so the database is current.

You need to choose the size and number of buffers that give you the maximum performance benefit. If your database uses OSAM, you might also decide to use OSAM sequential buffering. The subtopics in this topic can help you with these decisions.

Multiple Buffers in Virtual Storage

You can specify both the number of buffers needed in virtual storage and their size. You can specify multiple buffers with different sizes. Because a complete block or CI is read into a buffer, the buffer must be at least as large as the block or CI that is read into it. For best performance, use multiple buffers in virtual storage. To understand why, you need to understand the concept of buffers and how they are used in virtual storage.

When the data an application program needs is already in a buffer, the data can be used immediately. The application program is not forced to wait for the data to be read from the database to the buffer. Because the application program does not wait, performance is better. By having multiple buffers in virtual storage and by making a buffer large enough to contain all the segments of a CI or block, you increase the chance that the data needed by application programs is already in virtual storage. Thus, the reason for having multiple buffers in virtual storage is to eliminate some of an application program's wait time.

In virtual storage, all buffers are put in a buffer pool. Separate buffer pools exist for VSAM and OSAM. A buffer pool is divided into subpools. Each subpool is defined with a subpool definition statement. Each subpool consists of a specified number of buffers of the same size. With OSAM and VSAM you can specify multiple subpools with buffers of the same size.

"Use" Chain

In the subpool, buffers are chained together in the order in which they have been used. This organization is called a "use chain." The most recently used buffers are at the top of the use chain and the least recently used buffers are at the bottom.

The Buffer Handler

When a buffer is needed, an internal component called the buffer handler selects the buffer at the bottom of the use chain, because buffers that are least recently used are less likely to contain data an application program needs to use again. If a selected buffer contains data an application program has modified, the contents of the buffer are written back to the database before the buffer is used. This causes the application program wait time discussed earlier.

Buffering Options

Background Write Option

If you use VSAM, you can reduce or eliminate wait time by using the background write option. This option is discussed under “VSAM Options” on page 264. Otherwise, you control and reduce wait time by carefully choosing of the number and size of buffers.

Shared Resource Pools

You can define multiple VSAM local shared resource pools. Multiple local shared resource pools allow you to specify multiple VSAM subpools of the same size. You create multiple shared resource pools and then place in each one a VSAM subpool that is the same size as other VSAM subpools in other local shared resource pools. You can then assign a specific database data set to a specific subpool by assigning the data set to a shared resource pool. The data set is directed to a specific subpool within the assigned shared resource pool based on the data set’s control interval size.

Using Separate Subpools

If you have many VSAM data sets with similar or equal control interval sizes, you might get a performance advantage by replacing a single large subpool with separate subpools of identically sized buffers. Creating separate subpools of the same size for VSAM data sets offers benefits similar to OSAM multiple subpool support.

You can also create separate subpools for VSAM KSDS index and data components within a VSAM local shared resource pool. Creating separate subpools can be advantageous because index and data components do not need to share buffers or compete for buffers in the same subpool.

Hiperspace Buffering

Multiple VSAM local shared resource pools enhance the benefits provided by Hiperspace™ buffering. Hiperspace buffering allows you to extend the buffering of 4K and multiples of 4K buffers to include buffers allocated in expanded storage in addition to the buffers allocated in virtual storage. Using multiple local shared resource pools and Hiperspace buffering allows data sets with certain reference patterns (for example, a primary index data set) to be isolated to a subpool backed by Hiperspace, which reduces the VSAM read I/O activity needed for database processing.

Hiperspace buffering is activated at IMS initialization. In batch systems, you place the necessary control statements in the DFSVSAMP data set. In online systems, you place the control statements in the IMS.PROCLIB data set with the member name DFSVSMnn. Hiperspace buffering is specified for VSAM buffers through one or two optional parameters applied to the VSRBF subpool definition statement.

Related Reading: For a brief explanation of how to specify hiperspace buffering, see “Hiperspace Buffering Parameters” on page 412.

Buffer Size

Pick buffer sizes that are equal to or larger than the size of the CIs and blocks that are read into the buffer. A variety of valid buffer sizes exist. If you pick buffers larger than your CI or block sizes, virtual storage is wasted.

For example, suppose your CI size is 1536 bytes. The smallest valid buffer size that can hold your CI is 2048 bytes. This wastes 512 bytes (2048 - 1536) and is not a good choice of CI and buffer size.

Buffer Numbers

Pick an appropriate number of buffers of each size so buffers are available for use when they are needed, an optimum amount of data is kept in virtual storage during application program processing, and application program wait time is minimized. The trade-off in picking a number of buffers is that each buffer uses up virtual storage.

When you initially choose buffer sizes and the number of buffers, you are making a scientific guess based on what you know about the design of your database and the processing requirements of your applications. After you choose and implement buffer size and numbers, various monitoring tools are available to help you determine how well your scientific guess worked. Monitoring is discussed in Chapter 14, "Monitoring Databases," on page 341.

Buffer size and number of buffers are specified when the system is initialized. Both can be changed (tuned) for optimum performance at any time. Tuning is discussed in Chapter 15, "Tuning Databases," on page 347.

VSAM Buffer Sizes

The buffer sizes (in bytes) that you can choose when using VSAM as the access method are:

512
1024
2048
4096
8192
12288
16384
20480
24576
28672
32768

In order not to waste buffer space, choose a buffer size that is the same as a valid CI size. Valid CI sizes for VSAM data clusters are:

- For data components up to 8192 bytes (or 8K bytes), the CI size must be a multiple of 512.
- For data components over 8192 bytes (or 8K bytes), the CI size must be a multiple of 2048 (up to a maximum of 32768 bytes).

Valid CI sizes (in bytes) for VSAM index clusters using VSAM catalogs are:

512
1024
2048
4096

Valid CI sizes for VSAM index clusters using integrated catalog facility catalogs are:

Buffering Options

- For index components up to 8192 bytes (or 8K bytes), the CI size must be a multiple of 512.
- For index components over 8192 bytes (or 8K bytes), the CI size must be a multiple of 2048 (up to a maximum of 32768 bytes).

OSAM Buffer Sizes

The buffer sizes (in bytes) that you can choose when using OSAM as the access method are:

512

1024

2048

Any multiple of 2048 up to a maximum of 32768

For OSAM data sets, choose a buffer size that is the same as a valid block size so that buffer space is not wasted. Valid block sizes for OSAM data sets are any size from 18 to 32768 bytes.

Restriction: When using sequential buffering and the coupling facility for OSAM data caching, the OSAM database block size must be defined in multiples of 256 bytes (decimal). Failure to define the block size accordingly can result in ABENDS0DB from the coupling facility. This condition exists even if the IMS system is accessing the database in read-only mode.

Specifying Buffers

Specify the number of buffers and their size when the system is initialized. Your specifications, which are given to the system in the form of control statements, are put in the:

- DFSVSAMP data set in batch, utility.
- IMS.PROCLIB data set with the member name DFSVSMnn in IMS DCCTL and DBCTL environments.

The following example shows the necessary control statements specifications:

- Four 2048-byte buffers for OSAM
- Four 2048-byte buffers and fifteen 1024-byte buffers for VSAM

```
//DFSVSAMP DD *  
  
:  
VSRBF=2048,4  
VSRBF=1024,15  
IOBF=(2048,4)  
/*
```

Detailed information on how to code these control statements is located in the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

OSAM buffers can be fixed in storage using the IOBF= parameter. In VSAM, buffers are fixed using the VSAMFIX= parameter in the OPTIONS statement. This parameter is described under “VSAM Options” on page 264. Performance is generally improved if buffers are fixed in storage, then page faults do not occur. A page fault occurs when an instruction needs a page (a specific piece of storage) and the page is not in storage.

With OSAM, you can fix the buffers and their buffer prefixes, or the buffer prefixes and the subpool header, in storage. In addition, you can selectively fix buffer subpools, that is, you can choose to fix some buffer subpools and not others. Buffer subpools are fixed using the IOBF= parameter. The format of this parameter is:

IOBF= (*length,number,fix1,fix2,id*)

where:

- *length* is the size of buffers in a subpool.
- *number* is the number of buffers in a subpool. If three or fewer are specified, IMS gives you three; otherwise, it gives you the number specified. If you do not specify a sufficient number of buffers, your application program calls could waste time waiting for buffer space.
- *fix1* is whether the buffers and buffer prefixes in this subpool need to be fixed and is specified as Y or N (yes or no).
- *fix2* is whether the buffer prefixes in this subpool and the subpool header need to be fixed and is specified as Y or N (yes or no).

The default for the fix1 parameter is that buffers and their prefixes are not fixed. The default for the fix2 parameter is that buffer prefixes and the subpool header are not fixed.

- *id* is a parameter that specifies an identifier to be assigned to the subpool. It is used in conjunction with the DBD statement to assign a specific subpool to a given data set. This DBD statement is not the DBD statement used in a DBD generation but one specified during execution, as described in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*. The id parameter allows you to have more than one subpool with the same buffer size. You can use it to:
 - Get better distribution of activity among subpools
 - Direct new database applications to “private” subpools
 - Control the contention between a BMP and MPPs for subpools

OSAM Sequential Buffering

Sequential Buffering (SB) is an extension of the normal buffering technique used for OSAM database data sets. When SB is active, multiple consecutive blocks can be read from your database with a single I/O operation. (SB does not enhance OSAM write operations.) This technique can help reduce the elapsed time of many programs and utilities that sequentially process your databases.

Sequential Buffering Introduction

The normal OSAM buffering method reads only one block with each I/O operation. This method is known as a *random read*. Without SB, IMS must issue a random read each time your program processes a block that is not already in the OSAM buffer pool. For programs that process your databases sequentially, random reads can be time-consuming because the DASD must rotate one revolution or more between each read.

SB reduces the time needed for I/O read operations in three ways:

- By reading 10 consecutive blocks with a single I/O operation. This is called a *sequential read*. Sequential reads reduce the number of I/O operations necessary to sequentially process a database data set.

When a sequential read is issued, the block containing the segment your program requested plus nine adjacent blocks are read from the database into an

OSAM Sequential Buffering

SB buffer pool in virtual storage. When your program processes segments in any of the other nine blocks, no I/O operations are required because the blocks are already in the SB buffer pool.

Example: If your program sequentially processes an OSAM data set containing 100,000 consecutive blocks, 100,000 I/O operations are required using the normal OSAM buffering method. SB can take as few as 10,000 I/O operations to process the same data set.

- By monitoring the database I/O reference pattern and deciding if it is more efficient to satisfy a particular I/O request with a sequential read or a random read. This decision is made for each I/O request processed by SB.
- By overlapping sequential read I/O operations with CPC processing and other I/O operations of the same application. When overlapped sequential reads are used, SB anticipates future requests for blocks and reads those blocks into SB buffers before they are actually needed by your application. (Overlapped I/O is supported only for batch and BMP regions.)

Benefits of Sequential Buffering

By using SB, any application program or utility that sequentially processes OSAM data sets can run faster. Because many other factors affect the elapsed time of a job, the time savings is difficult to predict. You need to experiment with SB to determine actual time savings.

Programs That Can Benefit from SB

Some of the programs, utilities, and functions that might benefit from the use of SB are:

- IMS batch programs that sequentially process your databases.
- BMPs that sequentially process your databases.
- Any long-running MPP, Fast Path, and CICS programs that sequentially process your databases.

Note: SB is possible but not recommended for short-running MPP, IFP, and CICS programs. SB is not recommended for the short-running programs, because SB has a high initialization overhead each time such online programs are run.

- IMS utilities, including:
 - Online Database Image Copy
 - HD Reorganization Unload
 - Partial Database Reorganization
 - Surveyor
 - Database Scan
 - Database Prefix Update
 - Batch Backout
- HALDB Online Reorganization function

Typical Productivity Benefits of SB

By using SB for programs and utilities that sequentially process your databases, you might be able to:

- Run existing sequential application programs within decreasing “batch window times.” For example, if the time you set aside to run batch application programs is reduced by one hour, you might still be able to run all the programs you normally run within this reduced time period.

- Run additional sequential application programs within the same time period.
- Run some sequential application programs more often.
- Make online image copies much faster.
- Reduce the time needed to reorganize your databases.

Flexibility of SB Use

IMS provides several methods for requesting SB. You can request the use of SB for specific programs and utilities during PSBGEN or by using SB control statements. You can also request the use of SB for all or some batch and BMP programs by using an SB Initialization Exit Routine.

IMS also allows a system programmer or master terminal operator (MTO) to override requests for the use of SB by disallowing its use. This is done by issuing an SB MTO command or using an SB Initialization Exit Routine. The use of SB can be disallowed during certain times of the day to avoid virtual or real storage constraint problems.

These methods of controlling the use of SB are discussed in “How to Request the Use of SB” on page 261.

How SB Buffers Data

This topic describes what happens when you request SB. You will learn what SB buffers, how and when SB is activated, and what happens to the data that SB buffers.

What SB Buffers

As discussed in Chapter 8, “Choosing Optional Database Functions,” on page 157, HD databases can consist of multiple data set groups. A database PCB can therefore refer to several data set groups. A database PCB can also refer to several data set groups when the database referenced by the PCB is involved in logical relationships. A particular database, and therefore a particular data set group, can be referenced by multiple database PCBs. A specific data set group referenced by a specific database PCB is referred to in the following discussion as a *DB-PCB/DSG pair*.

When SB is activated, it buffers data from the OSAM data set associated with a specific DB-PCB/DSG pair. SB can be active for several DB-PCB/DSG pairs at the same time, but each pair requires a separate activation.

Conditional Activation and Periodical Evaluation of SB

IMS does not immediately activate SB when you request it. Instead, when SB is requested for a program, IMS begins monitoring the I/O reference pattern and activity rate for each DB-PCB/DSG pair used by the program. After awhile, IMS performs the first of a series of *periodical evaluations* of the buffering process. IMS performs these periodic evaluation for each DB-PCB/DSB pair. This periodical evaluation determines if the use of SB would be beneficial for the DB-PCB/DSG pair. If the use of SB would be beneficial, IMS activates SB for the DB-PCB/DSG pair. This activation of SB is known as *conditional activation*.

After SB is activated, IMS continues to periodically evaluate the I/O reference pattern and activity rate. Based on these evaluations, IMS can:

- Temporarily deactivate SB and continue to monitor the I/O reference pattern and activity rate. Temporary deactivation is implemented to unfix and page-release the SB buffers.

OSAM Sequential Buffering

- Temporarily deactivate monitoring of the I/O reference pattern and activity rate. This form of temporary deactivation is implemented only if SB has been deactivated and IMS concludes from subsequent evaluations that use of SB would still not be beneficial.

When SB is temporarily deactivated, it can be reactivated later based on the results of subsequent evaluations.

Individual periodical evaluations are performed for each DB-PCB/DSG pair. Therefore, IMS can deactivate SB for one DB-PCB/DSG pair while SB remains active for other DB-PCB/DSG pairs.

Role of the SB Buffer Handler

When SB is activated for a DB-PCB/DSG pair, a pool of SB buffers is allocated to the pair (SB buffers are also discussed in “Virtual Storage Considerations for SB”). Each SB buffer pool consists of n buffer sets (the default is four) and each buffer set contains 10 buffers. These buffers are used by an internal component called the SB buffer handler to hold the sets of 10 consecutive blocks read with sequential reads.

While SB is active, all requests for database blocks not found in the OSAM buffer pool are sent to the SB buffer handler. The SB buffer handler responds to these requests in the following way:

- If the requested block is already in an SB buffer, a copy of the block is put into an OSAM buffer.
- If the requested block is not in an SB buffer, the SB buffer handler analyzes a record of previous I/O requests and decides whether to issue a sequential read or a random read. If it decides to issue a random read, the requested block is read directly into an OSAM buffer. If it decides to issue a sequential read, the requested block and nine adjacent blocks are read into an SB buffer set. When the sequential read is complete, a copy of the requested block is put into an OSAM buffer.
- The SB buffer handler also decides when to initiate overlapped sequential reads.

Note: When processing a request from an online program, the SB buffer handler only searches the SB buffer pools allocated to that online program.

Related Reading: For information on how IMS invalidates SB buffers, see the data-sharing chapter of *IMS Version 9: Administration Guide: System*.

Virtual Storage Considerations for SB

Each DB-PCB/DSG pair buffered by SB has its own SB buffer pool. By default, each SB buffer pool contains four buffer sets (although IMS lets you change this value). Ten buffers exist in each buffer set. Each buffer is large enough to hold one OSAM data set block.

The total size of each SB buffer pool is:

$$4 * 10 * \text{block size}$$

The SB buffers are page-fixed in storage to eliminate page faults, reduce the path length of I/O operations, and increase performance. SB buffers are page-unfixed and page-released when a periodical evaluation temporarily deactivates SB.

You must ensure that the batch, online or DBCTL region has enough virtual storage to accommodate the SB buffer pools. This storage requirement can be considerable, depending upon the block size and the number of programs using SB.

SB is not recommended in real storage-constrained environments such as batch and DB/TM.

Some systems are storage-constrained only during certain periods of time, such as during online peak times. You can use an SB Initialization Exit Routine to control the use of SB according to specific criteria (the time) of day.

Related Reading: For details on the SB Initialization User Exit Routine see *IMS Version 9: Customization Guide*.

How to Request the Use of SB

IMS provides two methods for specifying which of your programs and databases should use SB.

1. You can explicitly specify which programs and utilities should use SB during PSB generation or by using SB control statements.
2. You can specify that by default all or a subset of your batch and BMP programs and utilities should use SB by coding an SB exit routine or by using a sample SB exit routine provided with IMS.

Determine which method you will use. Using the second method is easier because you do not need to know which BMP and batch programs use sequential processing. However, using SB by default can lead to an uncontrolled increase in real and virtual storage use, which can impact system performance. Generally, if you are running IMS in a storage-constrained z/OS environment, use the first method. If you are running IMS in a non storage-constrained z/OS environment, use the second method.

Requesting SB During PSB Generation

To request SB during PSB generation, specify SB=COND in the PCB macro instruction of your application's PSB. (This is not possible for IMS utilities that do not use a PSB during execution.) You code this keyword for each database PCB buffered with SB.

The following diagram shows the syntax of the SB keyword in the PCB statement.



COND Specifies that SB should be conditionally activated for this PCB.

NO Specifies that SB should not be used for this PCB.

If you do not include the SB keyword in your PCB, IMS defaults to NO unless specified otherwise in the SB exit routine.

The SB keyword value can be overridden by SB control statements. This option is discussed in “Requesting SB With SB Control Statements” on page 262.

The following example shows a PCB statement coded to request conditional activation of SB:

OSAM Sequential Buffering

```
SKILLA   PCB   TYPE=DB,DBDNAME=SKILLDB,KEYLEN=100,  
          PROCOPT=GR,SB=COND
```

Detailed instructions for coding PSB statements are contained in *IMS Version 9: Utilities Reference: System*.

Requesting SB With SB Control Statements

You can put SBPARM control statements in the optional //DFSCTL file. This file is defined by a //DFSCTL DD statement in the JCL of your batch, dependent, or online region. You can use the SBPARM control statement to:

- Specify which database PCBs (and which data sets referenced by the database PCB) should use SB
- Override the default number of buffer sets

This control statement allows you to override PSB specifications without requiring you to regenerate the PSB.

You can specify keywords that request use of SB for all or specific DBD names, DD names, PSB names, and PCB labels. You can also combine these keywords to further restrict when SB is used.

By using the BUFSETS keyword of the SBPARM control statement, you can change the number of buffer sets allocated to SB buffer pools. For details on the SB buffer pools see “Virtual Storage Considerations for SB” on page 260. The default number of buffer sets is four. Badly organized databases can require six or more buffer sets for efficient sequential processing. Well-organized databases require as few as two buffer sets. An indicator of how well-organized your database is can be found in the optional //DFSSTAT reports.

Related Reading:

- For details on //DFSSTAT reports, see *IMS Version 9: Utilities Reference: Database and Transaction Manager*.
- For information on tuning the number of buffer sets, see Chapter 15, “Tuning Databases,” on page 347.

The example below shows the SBPARM control statement necessary to request conditional activation of SB for all DBD names, DD names, PSB names, and PCBs.

```
SBPARM ACTIV=COND
```

The next example shows the parameters necessary to:

- Request conditional activation of SB for all PCBs that were coded with 'DBDNAME=SKILLDB' during PSB generation
- Set the number of buffer sets to 6

```
SBPARM ACTIV=COND,DB=SKILLDB,BUFSETS=6
```

Detailed instructions for coding the SBPARM control statement are contained in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Requesting SB with an SB Initialization Exit Routine

You can use an SB Initialization Exit Routine to:

- Request conditional activation of SB for all or some batch and BMP programs
- Allow or disallow the use of SB
- Change the default number of buffer sets

You can do this by writing your own SB exit routine or by selecting a sample SB exit routine and copying it under the name DFSSBUX0 into IMS.SDFSRESL. An SB exit routine allows you to dynamically control the use of SB at application scheduling time.

IMS supplies five sample SB exit routines in IMS.SDFSSRC and IMS.SDFSRESL. Three of the sample routines request SB for various subsets of application programs and utilities. One sample routine requests SB during certain times of the day and another routine disallows use of SB. You can use these sample routines as written or modify them to fit your needs.

Detailed instructions for the SB Initialization Exit Routine are in the *IMS Version 9: Customization Guide*.

SB Options or Parameters Provided by Several Sources

If you provide the same SB option or parameter in more than one place, the following priority list applies (item 1 having the highest priority):

1. SB control statement specifications (the n th control statement overrides the m th control statement, where $n > m$)
2. PSB specifications
3. Defaults changed by the SB Initialization Exit Routine
4. IMS defaults

Using SB in an Online System

To allow the use of SB in an online IMS or DBCTL environment, an IMS system programmer must explicitly request that IMS load the SB modules. This is done by putting an SBONLINE control statement in the DFSVSMxx member. By default, IMS does not load SB modules in an online environment. This helps avoid a noticeable increase in virtual storage requirements.

The two forms of the SBONLINE control statement are:

SBONLINE

or

SBONLINE,MAXSB=nnnnn

where *nnnnn* is the maximum storage (in kilobytes) that can be used for SB buffers.

When the MAXSB limit is reached, IMS stops allocating SB buffers to online applications until terminating online programs release SB buffer space. By default, if you do not specify the MAXSB= keyword, the maximum storage for SB buffers is unlimited.

Detailed instructions for coding the SBONLINE control statement are contained in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Disallowing the Use of SB

This topic describes how an IMS system programmer or MTO can disallow the use of SB. When the use of SB has been disallowed, a request for conditional activation of SB is ignored.

There are three ways to disallow the use of SB. The following list describes the three methods:

OSAM Sequential Buffering

1. An SB Initialization Exit Routine can be written (or a sample exit routine adapted) that can dynamically disallow and allow use of SB. This method can be used if you are using SB in an IMS batch, online, or DBCTL environment.
2. The MTO commands /STOP SB and /START SB can be issued to dynamically disallow and allow use of SB within an IMS online subsystem.

Related Reading: For details on the /STOP SB and /START SB commands, see *IMS Version 9: Command Reference*.

3. The SBONLINE control statement can be omitted from the DFSVSMxx member. This will keep IMS from loading the SB modules into the online subsystem. No program in the online subsystem will be able to use SB.

VSAM Options

Several types of options can be chosen for databases using VSAM. Specifying options such as free space for the ESDS data set, logical record size, and CI size are discussed in the preceding topics in this chapter. This topic describes these optional functions:

1. Functions specified in the OPTIONS control statement when IMS is initialized.
2. Functions specified in the POOLID, VSRBF, and DBD control statements when IMS is initialized.
3. Functions specified in the Access Method Services DEFINE CLUSTER command when a data set is defined.

Optional Functions Specified in the OPTIONS Control Statement

Several options exist that can be chosen during IMS system initialization for databases using VSAM. These options are specified in the OPTIONS control statement. In a batch system, the options you specify are put in the data set with the DDNAME DFSVSAMP. In an online system, they are put in the IMS.PROCLIB data set with the member name DFSVSMnn. Your choice of VSAM options can affect performance, use of space in the database, and recovery. This topic describes each option and the implications of using it.

The OPTIONS statement is described in detail in the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*. The OPTIONS statement and all its parameters are optional.

Using Background Write (BGWRT Parameter)

When an application program issues a call requiring that data be read from the database, the data is read into a buffer. If the buffer the data is to be read into contains altered data, the altered data must be written back to the database before the buffer can be used. If the data was not written back to the database, the data would be lost (overlaid) when new data was read into the buffer. Then there would be no way to update the database.

For these reasons, when an application program needs data read into a buffer and the buffer contains altered data, the application program waits while the buffer is written to the database. This waiting time decreases performance. The application program is ready to do processing, but the buffer is not available for use. Background write is a function you can choose in the OPTIONS statement that reduces the amount of wait time lost for this reason.

To understand how background write works, you need to know something about how buffers are used in a subpool. You specify the number of buffers and their size. All buffers of the same size are in the same subpool. Buffers in a subpool are on a

use chain, that is, they are chained together in the order in which they have been most or least recently used. The most recently used buffers are at the top of the use chain; least recently used buffers are at the bottom.

When a buffer is needed, the VSAM buffer manager selects the buffer at the bottom of the use chain. The buffer at the bottom of the use chain is selected, because buffers that have not been used recently are less likely to contain data that will be used again. If the buffer the VSAM buffer handler picks contains altered data, the data is written to the database before the buffer is used. It is during this step that the application program is waiting.

Background write solves the following problem: when the VSAM buffer manager gets a buffer in any subpool, it looks (when background write is used) at the next buffer on the use chain. The next buffer on the use chain will be used next. If the buffer contains altered data, IMS is notified so background write will be invoked. Background write has VSAM write data to the database from some percentage of the buffers at the bottom of the use chain. VSAM does this for *all* subpools. The data that is written to the database still remains in the buffers so the application program can still use any data in the buffers.

Background write is a very useful function when processing is done sequentially, but it is not as important to use in online systems as in batch. This is because, in online environments, IMS automatically writes buffers to the database at sync points.

To use background write, specify BGWRT=YES,n on the OPTIONS statement, where n is the percentage of buffers in each subpool to be written to the database. If you do not code the BGWRT= parameter, the default is BGWRT=YES and the default percentage is 34%. If an application program continually uses buffers but does not reexamine the data in them, you can make n 99%. Then, a buffer will normally be available when it is needed.

CICS does not support this function.

Choosing an Insert Strategy (INSERT Parameter)

Get free space in a CI in a KSDS is by specifying it in the DEFINE CLUSTER command. (The DEFINE CLUSTER command is explained in “Specifying Free Space for a KSDS (FREESPACE Parameter)” on page 267. Free space for a KSDS cannot be specified using the FRSPC= keyword in the DBD.

To specify free space in the DEFINE CLUSTER command, you must decide:

- Whether free space you have specified is preserved or used when more than one root segment is inserted at the same time into the KSDS.
- Whether to split the CI at the point where the root is inserted, or midway in the CI, when a root that causes a CI split is inserted.

These choices are specified in the INSERT= parameter in the OPTIONS statement. INSERT=SEQ preserves the free space and splits the CI at the point where the root is inserted. INSERT=SKP does not preserve the free space and splits the CI midway in the CI. In most cases, specify INSERT=SEQ so free space will be available in the future when you insert root segments. Your application determines which choice gives the best performance.

If you do not specify the INSERT= parameter, the default is INSERT=SKP.

Using the IMS Trace Parameters

The IMS trace parameters trace information that has proven valuable in solving problems in the specific area of the trace. All traces share sequencing numbers so that a general picture of the IMS environment can be obtained by looking at all the traces.

ON is the default for the IMS DL/I, LOCK and retrieve traces. OFF is the default for all other traces. The traces can be turned on at IMS initialization time. They can also be started or stopped by the /TRACE command during IMS execution. Output from long-running traces can be saved on the system log if requested.

Related Reading: For more information on the trace parameters, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Determining Which Dump Option to Use (DUMP Parameter)

The dump option is a serviceability aid that has no impact on performance. It merely describes the type of abend to take place if an abend occurs in the buffer handler (an internal component). If DUMP=YES is specified, the control region will abend when there is an abend in the buffer handler.

Deciding Whether to Fix VSAM Database Buffers and IOBs in Storage (VSAMFIX Parameter)

Each VSAM subpool contains buffers and input/output control blocks (IOBs). Performance is generally improved if these buffers and IOBs are fixed in storage. Then, page faults do not occur. A page fault occurs when an instruction references a page (a specific piece of storage) that is not in real storage.

You can specify whether buffers and IOBs are fixed in storage in the VSAMFIX= parameter of the OPTIONS statement. If you have buffers or IOBs fixed, they are fixed in all subpools. If you do not code the VSAMFIX= parameter, the default is that buffers and IOBs are not fixed.

This parameter can be used in a CICS environment if the buffers were specified by IMS.

Using Local Shared Resources (VSAMPLS Parameter)

Specifying VSAMPLS=LOCL in the OPTIONS statement is for local shared resources (LSR). When you specify VSAMPLS=LOCL, VSAM control blocks and subpools are put in the IMS control region. VSAMPLS=LOCL is the only valid operand and the default.

Optional Functions Specified in the POOLID, DBD, and VSRBF Control Statements

Options chosen during IMS initialization determine the size and structure of VSAM local shared resource pools. In a batch environment, you specify these options in a data set with the DDNAME DFSVVSAMP. In online systems, you specify these options in the IMS.PROCLIB data set with the member name DFSVSMnn.

With these options, you can enhance IMS performance by:

- Defining multiple local shared resource pools
- Dedicating subpools to a specific data set
- Defining separate subpools for index and data components of VSAM data sets

Related Reading: Implementing the POOLID, VSRBF, and DBD control statements and their corresponding parameters is described in detail in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Optional Functions Specified in the Access Method Services DEFINE CLUSTER Command

There are several optional functions that affect performance that can be chosen when you define your VSAM data sets. These functions are specified in the Access Method Services DEFINE CLUSTER command. HALDBs require that the REUSE parameter be specified on the DEFINE CLUSTER command. IMS Online Recovery Services takes advantage of the REUSE parameter, if it is specified.

Related Reading: This command and all its parameters are described in detail in *DFSMS Access Method Services for Catalogs*.

Specifying that 'Fuzzy' Image Copies Can be Taken with the Database Image Copy 2 (DFSUDMT0)

To establish that 'fuzzy' image copies of KSDSs can be taken with the Database Image Copy 2 (DFSUDMT0), specify the BWO(TYPEIMS) parameter. For this option to take effect the following conditions must exist:

- The KSDS must be SMS-managed.
- All access to the KSDS, once this option is specified, is done under DFSMS 1.3 or later version (once the KSDS has been opened under DFSMS 1.3, attempts to open it under an earlier version will fail).

Specifying Free Space for a KSDS (FREESPACE Parameter)

If get free space in a CI in a KSDS, specify it in the FREESPACE parameter in the DEFINE CLUSTER command. Free space for a KSDS can not be specified using the FRSPC= keyword in the DBD.

You specify free space in the FREESPACE parameter as a percentage. The format of the parameter is FREESPACE(x,y) where:

- x** is the percentage of space in a CI left free when the database is loaded or when a CI split occurs after initial load
- y** is the percentage of space in a control area (CA) left free when the database is loaded or when a CA split occurs after initial load.

Free space is preserved when a CI or CA is split by coding INSERT=SEQ in the OPTIONS control statement. INSERT=SEQ is explained in "Choosing an Insert Strategy (INSERT Parameter)" on page 265.

If you do not specify the FREESPACE parameter, the default is that no free space is reserved in the KSDS data set when the database is loaded.

Specifying Whether Data Set Space Is Pre-formatted for Initial Load (SPEED | RECOVERY Parameter)

When initially loading a VSAM data set, you can specify whether you need the data set pre-formatted in the SPEED | RECOVERY parameter. When SPEED is specified, it says the data set should *not* be pre-formatted. An advantage of pre-formatting a data set is; if initial load fails, you can recover and continue loading database records after the last correctly-written record. However, IMS does *not* support the RECOVERY option (except by use of the Utility Control Facility). So, although you can specify it, you cannot perform recovery. Because you cannot take

advantage of recovery when you specify the RECOVERY parameter, you should specify SPEED to improve performance during initial load.

To be able to recover your data set during load, you should load it under control of the Utility Control Facility. This utility is described in *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

RECOVERY is the default for this parameter.

Specifying Whether Index Set Records Are Replicated

A VSAM KSDS cluster has a data component (where segments are stored in HISAM, HIDAM, or PHIDAM databases) and an index component (called the VSAM index in this discussion.) The VSAM index contains pointers to CIs in the KSDS data component. When a specific key in a KSDS is requested, the VSAM index is used to limit the search for the CI that contains the correct root segment. Without the VSAM index, the entire KSDS data component could be searched to find the correct CI. The VSAM index can be on either the same volume as the data component or on another volume. It is the VSAM index whose options are of concern here. You need to know some things about the VSAM index before the options are described.

The VSAM index consists of one or more levels, as shown in Figure 171. The first (lowest) level is called the sequence set level. All other levels are called index set levels. The sequence set level has a sequence set record for each CA in the database. Each sequence set record contains a pointer to each CI in a specific CA and the highest root segment's key in that CI.

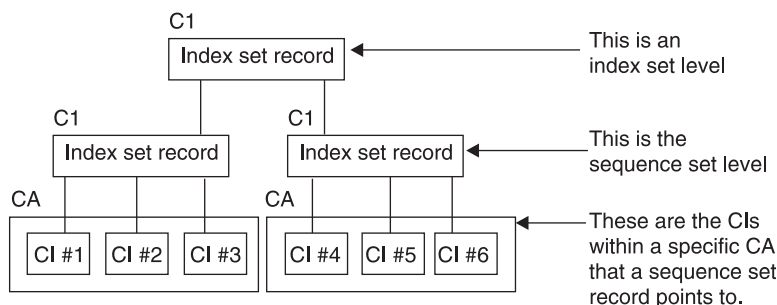


Figure 171. Levels in a VSAM Index

Index set records on the first index set level contain pointers to sequence set records. Each pointer on the first index set level contains the address of a sequence set record and the highest root segment key in the sequence set record pointed to.

If no more room exists for new pointers in an index set record, a new index set record is started on the same level. As soon as there are two index set records on a level, a new index set record is started on the next higher level.

At the second and higher levels of the index set, the pointers are to index set records at the next lowest level. Each pointer contains the address of an index set record at the next lower level along with the highest key in the index set record pointed to.

One option you can specify for the VSAM index that especially affects performance is the REPLICATE | NOREPLICATE parameter in the DEFINE CLUSTER command. If

you specify REPLICATE, each record in the sequence set and the index set is written as many times as it will fit on the track. Repeat records to reduce the delay caused when the disk rotates. The repetition of records means the arm is almost always close or over a record so very little disk rotation is necessary. Repeating records also improves performance. Note, however, that the VSAM index, because of the repetition, will probably require more direct-access space.

If you specify NOREPLICATE, records in the VSAM index are not repeated. NOREPLICATE is the default for this parameter.

There is a new option that you must specify for KSDSs in order to take 'fuzzy' image copies using the Database Image Copy 2 utility. BWO(TYPEIMS) is the specification. The KSDS must be SMS-managed for BWO(TYPEIMS) to mean anything. And, you should ensure that all access to the KSDS (once the BWO(TYPEIMS) option has been specified) is under DFSMS 1.3 or higher.

OSAM Options

Two types of options are available for databases using OSAM:

1. Options specified in the DBD (free space, logical record size, CI size).

These options are covered in preceding sections in this chapter.

2. Options specified in the OPTIONS control statement when IMS is initialized.

In a batch system, the options are put in the data set with the DDNAME DFSVSAMP. In an online system, they are put in the IMS.PROCLIB data set with the member name DFSVSMnn. Your choice of OSAM options can affect performance, recovery, and the use of space in the database.

The OPTIONS statement is described in detail in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*. The statement and all its parameters are optional.

Dump Option (DUMP Parameter)

The dump option is a serviceability aid that has no impact on performance. It merely describes the type of abnormal termination to take place when abnormal termination occurs in the buffer handler (an internal component).

Deciding Which FIELD Statements to Code in the DBD

Chapter 2, "Standards and Procedures," on page 19 describes the statements that are coded in the DBD. One of those statements is the FIELD statement, which defines a field within a segment type. An important thing to note about the FIELD statement is that it has to be coded for sequence fields and for fields an application program can refer to in the SSA of a call. A FIELD statement also has to be coded if it is referenced by a SENFLD statement in any PSB. Because each FIELD statement takes up storage in the DMB control block, do not generate FIELD statements that are unnecessary.

Planning for Maintenance

In designing your database, remember to plan for maintenance. If your applications require, for instance, that the database be available 16 hours a day, you do not design a database that takes 10 hours to unload and reload. No guideline we can give you for planning for maintenance exists, because all such plans are application dependent. However, remember to plan for it.

Planning for Maintenance

A possible solution to the problem just described is to make three separate databases and put them on different volumes. If the separate databases have different key ranges, then application programs could include logic to determine which database to process against. This solution would allow you to reorganize the three databases at separate times, eliminating the need for a single 10-hour reorganization. Another solution to the problem if your database uses HDAM or HIDAM might be to do a partial reorganization using the Partial Database Reorganization utility (described in Chapter 16, “Modifying Databases,” on page 429).

In the online environment, the Image Copy utilities allow you to do some maintenance without taking the database offline. These utilities let you take image copies of databases or partitions while they are allocated to and being used by an online IMS system.

HALDB provides greatly improved availability for large databases. By partitioning large databases, you can perform offline maintenance on a single partition, while the remaining partitions remain available.

I
I
I
I
You can also reorganize HALDBs online, which improves the performance of your HALDB without disrupting access to its data. If you plan to reorganize your HALDB online, make sure that there is enough DASD space to accommodate the reorganization process.

Related Reading: For information on reorganizing HALDBs online, see “HALDB Online Reorganization” on page 370.

Chapter 10. Designing Fast Path Databases

After you determine the type of database and optional functions that best suit your application's processing requirements, you need to make a series of decisions about database design and the use of options. This set of decisions primarily determines how well your database performs and how well it uses available space. These decisions are based on:

- The type of database and optional functions you have already chosen
- The performance requirements of your applications
- How much storage you have available for use online

This chapter examines the following topics:

- "Designing a Data Entry Database (DEDB)"
- "Designing a Main Storage Database (MSDB)" on page 277
- "High-Speed Sequential Processing (HSSP)" on page 283
- "Designing a DEDB or MSDB Buffer Pool" on page 286
- "Designing a DEDB Buffer Pool in the DBCTL Environment" on page 290

Designing a Data Entry Database (DEDB)

This topic describes the choices you need to make in designing a DEDB and proposes guidelines to help you make these choices.

To design a DEDB, you must know the following information:

- How the application fits the limitations imposed by the DEDB itself
- How the application can make optimum use of the area concept of a DEDB
- The size of the CI
- The size of the UOW
- The DEDB randomizing routine
- Record deactivation
- Multiple copies of an area data set
- PCL (physical child last pointer)
- Subset pointers

Related Reading: DEDBs can be shared. For information on DEDB data sharing, see *IMS Version 9: Administration Guide: System* and *IMS Version 9: Utilities Reference: System*.

DEDB Design Guidelines

The following list describes guidelines for designing DEDBs:

- Except for the relationship between a parent and its children, the logical structure (defined by the PCB) does not need to follow the hierarchic order of segment types defined by the DBD.

For example, SENSEG statements for DDEP segments can precede the SENSEG statement for the SDEP segment. This implementation prevents unqualified GN processing from retrieving *all* SDEP segments before accessing the first DDEP segments.

Designing a Data Entry Database

- Most of the time, SDEP segments are retrieved all at once, using the DEDB Sequential Dependent Scan utility. If you later must relate SDEP segments to their roots, you must plan for root identification as part of the SDEP segment data.
- A journal can be implemented by collecting data across transactions using a DEDB. To minimize contention, you should plan for an area with more than one root segment. For example, a root segment can be dedicated to a transaction/region or to each terminal. To further control resource contention, you should assign different CIs to these root segments, because the CI is the basic unit of DEDB allocation.
- Following is a condition you might be confronted with and a way you might resolve it. Assume that transactions against a DEDB record are recorded in a journal using SDEP segments and that a requirement exists to interrogate the last 20 or so of them.

SDEP segments have a fast insert capability, but on the average, one I/O operation is needed for each retrieved segment. The additional I/O operations could be avoided by inserting the journal data as both a SDEP segment and a DDEP segment and by limiting the twin chain of DDEP segments to 20 occurrences. The replace or insert calls for DDEP segments does not necessarily cause additional I/O, since they can fit in the root CI. The root CI is always accessed even if the only call to the database is an insert of an SDEP segment. The online retrieve requests for the journal items can then be responded to by the DDEP segments instead of the SDEP segments.

- As physical DDEP twin chains build up, I/O activity increases. The SDEP segment type can be of some help if the application allows it.

The design calls for DDEP segments of one type to be batched and inserted as a single segment whenever their number reaches a certain limit. An identifier helps differentiate them from the regular journal segments. This design prevents updates after the data has been converted into SDEP segments.

DEDB Area Design Guidelines

The following are some reasons why DEDBs are divided into areas and some related design considerations:

- DEDBs should be divided into areas in a way that makes sense for the application programs.

Example: A service bureau organization makes a set of applications available to its customers. The design calls for a common database to be used by all users of this set of applications. The area concept fits this design because the randomizing routine and record keys can be set so that data requests are directed to the user's area only. Furthermore, on the operational side, users can be given specific time slots. Their areas are allocated and deallocated dynamically without interrupting other services currently using the same DEDB.

National or international companies with business locations spanning multiple time zones might take advantage of the partitioned database concept. Because not all areas must be online all the time, data can be spread across areas by time zone.

Preferential treatment for specific records (specific accounts, specific clients, and so on.) can be implemented without using a new database, for example, by keeping more sequential dependent segments online for certain records. By putting together those records in one area, you can define a larger sequential dependent segment part and control the retention period accordingly.

- The impact of permanent I/O errors and severe errors can be reduced using a DEDB. DL/I requires that all database data sets, except for HALDBs, be available

all the time. With a DEDB, the data not available is limited only to the area affected by the failure. Because the DEDB utilities run at the level of the area, the recovery of the failing area can be done while the rest of the database is accessible to online processing. The currently allocated log volume must be freed by a /DBR AREA command and used in the recovery operation. Track recovery is also supported. The recovered area can then be dynamically allocated back to the operational environment.

Related Reading: Make multiple copies of DEDB area data sets to make data more available to application programs. See “Multiple Copies of an Area Data Set” on page 276.

- Space management parameters can vary from one area to another. This includes: CI size, UOW size, root addressable part, overflow part, and sequential dependent part. Also, the device type can vary from one area to the other.
- It is feasible to define an area on more than one volume and have one volume dedicated to the sequential dependent part. This implementation might save some seek time as sequential dependent segments are continuously added at the end of the sequential dependent part. The savings depends on the current size of the sequential dependent part and the blocking factor used for sequential dependent segments. If an area spans more than one volume, volumes must be of the same type.
- Only the independent overflow part of a DEDB is extendable. Sufficient space should be provided for all parts when DEDBs are designed. To extend the independent overflow part of a DEDB, you must follow the procedures in “Extending DEDB Independent Overflow Online” on page 464.

The /DISPLAY command and the POS call can help monitor the usage of auxiliary space. Unused space in the root addressable and independent overflow parts can be reclaimed through reorganization. It should be noted that, in the overflow area, space is not automatically reused by ISRT calls. To be reused at call time, the space must amount to an entire CI, which is then made available to the ISRT space management algorithm. Local out-of-space conditions can occur, although some available space exists in the database.

- Adding or removing an area from a DEDB requires a DBDGEN and an ACBGEN. Database reload is required if areas are added or deleted in the middle of existing areas. Areas added other than at the end changes the area sequence number assigned to the areas. The subsequent log records written reflect this number, which is then used for recovery purposes. If areas are added between existing areas, prior log records will be invalid. Therefore, an image copy must be made following the unload/reload. Be aware that the sequence of the AREA statements in the DBD determines the sequence of the MRMB entries passed on entry to the randomizing routine. An area does not need to be mounted if the processing does not require it, so a DBDGEN/ACBGEN is not necessary to logically remove an area from processing.
- Careful monitoring of the retention period of each log allows you to make an image copy of one area at a time. Also, because the High-Speed DEDB Direct Reorganization utility logs changes, you do not need to make an image copy following a reorganization.
- The area concept allows randomizing at the area level, instead of randomizing throughout the entire DEDB. This means the key might need to carry some information to direct the randomizing routine to a specific area.

Determining the Size of the CI

The choice of a CI size depends on the following factors:

Designing a Data Entry Database

- CI sizes of 512, 1 KB, 2 KB, 4 KB, and up to 28 KB in 4 KB increments are supported.
- Only one RAP exists per CI. The average record length has to be considered. In the base section of the root addressable part, a CI can be shared only by the roots that randomize to its RAP and their DDEP segments.
- Track utilization according to the device type.
- SDEP segment writes. A larger CI requires a fewer number of I/Os to write the same amount of SDEP segments.
- The maximum segment size, which is 28,552 bytes if using a 28 KB CI size.

Determining the Size of the UOW

The UOW is the unit of space allocation in which you specify the size of the root addressable and independent overflow parts.

Three factors might affect the size of the UOW:

1. The High-Speed DEDB Direct Reorganization utility (DBFUHDR0) runs on a UOW basis. Therefore, while the UOW is being reorganized, none of the CIs and data they contain are available to other processing.

A large UOW can cause resource contention, resulting in increased response time if the utility is run during the online period. A minor side effect of a large UOW is the space reserved on DASD for the “reorganization UOW,” which is no longer used, but retained for compatibility purposes.

A UOW that is too small can cause some overhead during reorganization as the utility switches from one UOW to the next with very little useful work each time. However, this might not matter so much if reorganization time is not critical.
2. The use of processing option P, (explained in “Processing Option P (PROCOPT=P)” on page 275). This consideration pertains to sequential processing using BMP regions. If the application program is coded to take advantage of the 'GC' status code, this status code must be returned frequently enough to fit in the planned sync interval.

Assume every root CI needs to be modified and that, for resource control reasons, each sync interval is allowed to process sequentially no more than 20 CIs of data. The size of the UOW should not be set to more than 20 CIs. Otherwise, the expected 'GC' status code would not be returned in time for the application program to trigger a sync point, release the resources, and not lose position in the database.

A UOW that is too small, such as the minimum of two CIs, can cause too many ‘unsuccessful database call’ conditions each time a UOW is crossed. On a 'GC' status code, no segment is returned and the call must be reissued after an optional SYNC or CHKP call.
3. The dependent overflow (DASD space) usage is more efficient with a large UOW than a small UOW.

See “SDEP CI Preallocation and Reporting” for a discussion of how the size of the UOW affects DEDB design.

SDEP CI Preallocation and Reporting

Because of data sharing, SDEP CIs cannot be allocated one at a time. Also, each data sharing system requires its own current CI. Therefore, a set of SDEP CIs are preallocated to each IMS on an allocation call. The number of CIs obtained by an IMS is a function of the system’s insert rate. The insert process obtains the current CI, not the area open process.

Because the insert process obtains the current CI, space use and reporting is complex. If a preallocation attempt cannot obtain the number of CIs requested, the ISRT or sync point call receives status FS, even if there is enough space for that particular call. The FS processing marks the area as full, and any subsequent smaller inserts also fail.

When there are few available SDEP CIs in an area, the number that can actually be used for SDEP inserts varies depending on the system's insert rate. Also, the command /DIS AREA calculates the number of SDEP CIs free as those available for preallocation and any unused CIs preallocated to the IMS issuing the command. Area close processing discards CIs preallocated to the IMS, and the unused CIs are lost until the SDEP Delete utility is run. Therefore, the number of unused CIs reported by the /DIS AREA command after area close processing is smaller because the preallocated CIs are no longer available.

Processing Option P (PROCOPT=P)

The PROCOPT=P option is specified during the PCB generation in the PCB statement or in the SENSEG statement for the root segment.

The option takes effect only if the region type is a BMP. If specified, it offers the following advantage:

Whenever an attempt is made to retrieve or insert a DEDB segment that causes a UOW boundary to be crossed, a 'GC' status code is set in the PCB but no segment is returned or inserted. The only calls for which this takes place are: G(H)U, G(H)N, POS, and ISRT.

Although crossing the UOW boundary has no particular significance for most applications, the 'GC' status code that is returned indicates this could be a convenient time to invoke sync point processing. This is because a UOW boundary is also a CI boundary. As explained for sequential processing, a CI boundary is a convenient place to request a sync point.

The sync point is invoked by either a SYNC or a CHKP call, but this normally causes position on all currently accessed databases to be lost. The application program then has to resume processing by reestablishing position first. This situation is not always easy to solve, particularly for unqualified G(H)N processing.

An additional advantage with this processing option is, if a SYNC or CHKP call is issued after a 'GC' status code, database position is kept. Database position is such that an unqualified G(H)N call issued after a 'GC' status code returns the first root segment of the next UOW. When a 'GC' status code is returned, no data is presented or inserted. Therefore, the application program should, optionally, request a sync point, reissue the database call that caused the 'GC' status code, and proceed. The application program can ignore the 'GC' status code, and the next database call will work as usual.

Database recovery and change accumulation processing must buffer all log records written between sync points. Sync points must be taken at frequent intervals to avoid exhausting available storage. If not, database recovery might not be possible.

DEDB Randomizing Routine Design

A DEDB randomizing module is required for placing root segments in a DEDB. The randomizing module is also required for retrieving root segments from a DEDB. One

Designing a Data Entry Database

or more such modules can be used with an IMS system. Only one randomizing module can be associated with each DEDB.

Related Reading: Refer to *IMS Version 9: Customization Guide* for register usage and a sample randomizing program exit (DBFHDC40).

The purpose of the randomizing module is the same as in HDAM processing. A root search argument key field value is supplied by the application program and converted into a relative root anchor point number. Because the entry and exit interfaces are different, DEDB and HDAM randomizing routines are not object code compatible. The main line randomizing logic of HDAM should not need modification if randomizing through the whole DEDB.

Some additional differences between DEDB and HDAM randomizing routines are as follows:

- The ISRT algorithm attempts to put the entire database record close to the root segment (with the exception of SDEP segments). No BYTES parameter exists to limit the size of the record portion to be inserted in the root addressable part.
- With the DEDB, only one RAP can be defined in each root addressable CI.
- CIs that are not randomized to are left empty.

Because of the area concept, some applications might decide to randomize in a particular area rather than through all the DEDB as in HDAM processing. Therefore, the expected output of such a randomizing module is made up of a relative root anchor point number in an area and the address of the control block (DMAC) representing the area selected.

Keys that randomize to the same RAP are chained in ascending key sequence.

DEDB logic runs in parallel, so DEDB randomizing routines must be reentrant. The randomizing routines operate out of the common storage area (CSA). If they use operating system services like LOAD, DELETE, GETMAIN, and FREEMAIN, the routines must abide by the same rules as described in *IMS Version 9: Customization Guide*.

Multiple Copies of an Area Data Set

The data in an area is in a VSAM data set called the area data set (ADS). Installations can create as many as seven copies (multiple area data sets, MADS) of each ADS, making the data more available to application programs.

Each copy of an ADS contains exactly the same user data. Fast Path maintains data integrity by keeping identical data in the copies during application processing. When an application program updates data in an area, Fast Path updates that data in each copy of the ADS. When an application program reads data from an area, Fast Path retrieves the requested data from any one of the available copies of the ADS. All copies of an ADS must have the same definition but can reside on different devices and on different device types. Using copies of ADS is also helpful in direct access device migration; for example, from a 3380 device to a 3390 device.

If an ADS fails to open during normal open processing of a DEDB, none of the copies of the ADS can be allocated, and the area is stopped. However, when open failure occurs during emergency restart, only the failed ADS is deallocated and stopped. The other copies of the ADS remain available for use.

Record Deactivation

If an error occurs while an application program is updating a DEDB, it is not necessary to stop the database or the area. IMS continues to allow application programs to access that area, and it only prevents them from accessing the control interval in error. If multiple copies of the ADS exist, one copy of the data is always available. (It is unlikely that the same control interval is in error in seven copies of the ADS.) IMS automatically deactivates a record when a count of 10 errors is reached.

Record deactivation minimizes the effect of database failures and errors to the data in these ways:

- If multiple copies of an area data set are used, and an error occurs while an application program is trying to update that area, the error does not need immediate correction. Other application programs can continue to access the data in that area through other available copies of that area.
- If a copy of an area has errors, you can create a new copy from existing copies of the ADS using the DEDB Data Set Create utility. The copy with the errors can then be destroyed.

Physical Child Last Pointers

The PCL pointer makes it possible to access the last physical child of a segment type directly from the physical parent. Using the INSERT rule LAST avoids the need to follow a potentially long physical child pointer chain.

Subset Pointers

Subset pointers help you avoid unproductive get calls when you need to access the last part of a long segment chain. These pointers divide a chain of segment occurrences under the same parent into two or more groups, or subsets. You can define as many as eight subset pointers for any segment type, dividing the chain into as many as nine subsets. Each subset pointer points to the start of a new subset.

Related Reading: For more information on defining and using subset pointers, see the topic about Processing DEDBs with Subset Pointers in *IMS Version 9: Application Programming: Database Manager*.

Restrictions: When you unload and reload a DEDB containing subset pointers, IMS does not automatically retain the position of the subset pointers. When unloading the DEDB, you must note the position of the subset pointers, storing the information in a permanent place. (For example, you could append a field to each segment, indicating which subset pointer, if any, points to that segment.) Or, if a segment in a twin chain can be uniquely identified, identify the segment a subset pointer is pointing to and add a temporary indication to the segment for reload. When reloading the DEDB, you must redefine the subset pointers, setting them to the segments to which they were previously set.

Designing a Main Storage Database (MSDB)

This topic describes the choices you might need to make in designing an MSDB and proposes guidelines to help you make these choices.

Consider the following list of questions when designing an MSDB database:

- How are virtual storage requirements for the database calculated?

Designing a Main Storage Database

- How are virtual storage requirements for the Fast Path buffer pool calculated?
- What are the storage requirements for the I/O area?
- Should FLD calls or other DL/I calls be used for improved MSDB and DEDB performance?
- How can the difference in resource allocation between an MSDB and a DL/I database be a key to good performance?
- What are the requirements in designing for minimum resource contention in a mixed-mode environment?
- How is the number of MSDB segments loaded into virtual storage controlled?
- What are the auxiliary storage requirements for an MSDB?
- How can an MSDB be checkpointed?

Calculating Virtual Storage Requirements for an MSDB

You can calculate the storage requirements for an MSDB as follows:

$$(L + 4)S + C + 14F + X$$

where:

S = the number of segments in the MSDB as specified by the member DBFMSDBx in the IMS.PROCLIB

L = the segment length as specified in the DBD member

C = 80 for non-related MSDBs without a terminal-related key, *or*
94 for the other types of MSDB

F = the number of fields defined in the DBD member

X = 2 if C + 14F is not a multiple of 4, OR
0 if C + 14F is a multiple of 4

MSDBs reside in the z/OS extended common storage area (ECSA).

Calculating Buffer Requirements

Details about calculating buffer requirements are in “Designing a DEDB or MSDB Buffer Pool” on page 286, along with other Fast Path buffer requirements. The following considerations apply during execution:

- Fast Path buffer requirements vary with the type of call to the MSDB.
- With a GHx/REPL call sequence, an entire segment is kept in the Fast Path buffer until a sync point is reached. If the total size of a series of segments exceeds the NBA (normal buffer allocation), the NBA parameter needs to be adjusted rather than using the OBA (overflow buffer) on a regular basis. You should accommodate the total number of segments used between sync points.
- When using a FLD call, the VERIFY and CHANGE logic reside in the Fast Path buffer.

Calculating the Storage for an Application I/O Area

A GHx/REPL call requires an I/O area large enough to accommodate the largest segment to be processed. The FLD call requires storage to accommodate the total field search argument (FSA) requirements.

Understanding Resource Allocation, a Key to Performance

The MSDB resource allocation scheme is different from that of DL/I. Since the MSDB is a key to good performance, it is important to understand it.

1. An MSDB record can be shared (S) by multiple users or be owned exclusively (E) by one user.
2. The same record can have both statuses (shared and exclusive) at the same time.
3. Updates to MSDBs are applied during sync point processing. The resource is always owned in exclusive mode for the duration of sync point processing.

The different enqueue levels of an MSDB record, when a record is enqueued, and the duration are summarized in Table 19.

Table 19. Levels of Enqueue of an MSDB Record

| Enqueue Level | When | Duration |
|---------------------|--|---|
| READ | GH with no update intent | VERIFY/get calls |
| | From call time until sync point (phase 1) ¹ | Call processing |
| HOLD | GH with no update intent | At sync point, to reapply VERIFYs |
| | From call time until sync point (phase 1) ¹ | Phase 1 of sync point processing, then released |
| UPDATE ² | At sync point, to apply the results of CHANGE, REPL, DLET, or ISRT calls | Sync point processing, then released |

Notes:

1. If there was no FLD/VERIFY call against this resource or if this resource is not going to be updated, it is released. Otherwise, if only FLD/VERIFY logic has to be reapplied, the MSDB record is enqueued at the HOLD level. If the same record is involved in an update operation, it is enqueued at the UPDATE level as shown in the table above.
2. At DLET/REPL call time, no enqueue activity takes place because it is the prior GH call that set up the enqueue level.

Table 20 shows that the status of an MSDB record depends on the enqueue level of each program involved. Therefore, it is possible for an MSDB record to be enqueued with the shared and exclusive statuses at the same time. For example, such a record can be shared between program A (GH call for update) and program B (GU call), but cannot be shared at the same time with a third program, C, which is entering sync point with update on the record.

Table 20. Example of MSDB Record Status: Shared (S) or Owned Exclusively (E)

| Enqueue Level in Program B | Enqueue Level in Program A | | |
|----------------------------|----------------------------|------|--------|
| | READ | HOLD | UPDATE |
| READ | S | S | E |
| HOLD | S | E | E |
| UPDATE | E | E | E |

Designing a Main Storage Database

The FLD/CHANGE call does not participate in any allocation; therefore, FLD/CHANGE calls can be executed even though the same database record is being updated during sync point processing.

If FLD/CHANGE and FLD/VERIFY calls are mixed in the same FLD call, when the first FLD/VERIFY call is encountered, the level of enqueue is set to READ for the remainder of the FLD call.

Designing to Minimize Resource Contention

One reason to use an MSDB is its fast access to data and high availability for processing. To maintain high availability, you should design to avoid the contention for resources that is likely to happen in a high transaction rate environment.

The following is a list of performance-related considerations. Some of the considerations do not apply exclusively to MSDBs, but they are listed to give a better understanding of the operational environment.

- Access by Fast Path transactions to DL/I databases and use of the alternate PCB should be kept to a minimum. Use of the alternate PCB should be kept to a minimum because FP transactions must contend for resources with IMS transactions (some of which could be long running). Also, common sync point processing is invoked and entirely serialized in the IMS control region.
- To avoid resource contention when sharing MSDBs between Fast Path and DL/I transactions, You should try to make commit processing often and to avoid long-running scans.
- GH for read/update delays any sync point processing that intends to update the same MSDB resource. Therefore, GH logic should be used only when you assume the referenced segments will not be altered until completion of the transaction. If the resource is being updated, release is at the completion of sync point. Otherwise, the release is at entry to sync point.
- The following consideration deals with deadlock prevention. Deadlock can occur if transactions attempt to acquire (GH calls) multiple MSDB resources.

Whenever a request for an MSDB resource exists that is already allocated and the levels involved are HOLD or UPDATE, control is passed to IMS to detect a potential deadlock situation. Increase in path length and response time results. The latter can be significant if a deadlock occurs, thus requiring the pseudo abend of the transaction.

In order to reduce the likelihood of deadlocks caused by resource contention, sync point processing enqueues (UPDATE level) MSDB resources in a defined sequence. This sequence is in ascending order of segment addresses. MSDB segments are acquired in ascending order of keys within ascending order of MSDB names, first the page-fixed ones then the pageable MSDBs.

The application programmer can eliminate potential deadlock situations at call time by also acquiring (GH calls) MSDB resources using the same sequence.

- From the resource allocation scheme discussed earlier, you probably realize that FLD logic should be used whenever possible instead of GH/REPL logic.
 - The FLD/VERIFY call results in an enqueue at the READ level, and if no other levels are involved, then control is not passed to IMS. This occurrence results in a shorter path length.
 - The FLD/CHANGE call, when not issued in connection with VERIFY logic does not result in any enqueue within either Fast Path or IMS.
 - FLD logic has a shorter path length through the Program Request Handler, since only one call to process exists instead of two needed for GH/REPL logic.

- The FLD/CHANGE call *never* waits for any resource, even if that same resource is being updated in sync point processing.
- The FLD/VERIFY call waits only for sync point processing during which the same resource is being updated.
- With FLD logic, the resource is held in exclusive mode only during sync point processing.

In summary, programming with FLD logic can contribute to higher transaction rates and shorter response times.

The following examples, Figure 172 and Figure 173, show how the MSDB record is held in exclusive mode:




Figure 172. First Example MSDB Record Held in Exclusive Mode

The following notes are for Figure 172:

1. MSDB record R1 is held in exclusive mode against:
 - Any MSDB calls except CHANGE calls
 - Any other sync point processing that intends to update the same record
2. MSDB record R1 is held in exclusive mode against:
 - Any other GH for update
 - Any other sync point processing that intends to update the same record



Figure 173. Second Example MSDB Record Held in Exclusive Mode

The following notes are for Figure 173.

1. MSDB record R1 is held in exclusive mode against:
 - Any MSDB calls except CHANGE calls
 - Any other sync point processing that intends to update the same record
2. MSDB record is held in exclusive mode for the duration of the FLD call against any other sync point processing that intends to update the same resource

Choosing MSDBs to Load and Page-Fix

Deciding which MSDBs to load and page-fix involves a trade-off between desired application performance and the amount of real storage available. This decision is made with total Fast Path application requirements in mind. IMS system initialization requires additional information before MSDBs can be loaded and page fixed. This information is specified in member DBFMSDBx of IMS.PROCLIB. This member is

Designing a Main Storage Database

called by executing the control region startup procedure IMS. The suffix 'x' matches the parameter supplied in the MSDB keyword of the EXEC statement in procedure IMS.

The control information that loads and page fixes MSDBs is in 80-character record format in member DBFMSDBx. Either you supply this information or it can be supplied by the output of the MSDB maintenance utility. When the /NRE command requests MSDBLOAD, the definition of the databases to be loaded is found in the DBFMSDBx procedure.

The definition in DBFMSDBx can represent a subset of the MSDBs currently on the sequential data set identified by DD statement MSDBINIT. Explicitly state each MSDB that you want IMS to load. If each MSDB is not explicitly stated, IMS abends.

The format for DBFMSDBx is as follows:

►►—DBD=*dbd_name*,—NBSEGS=*nnnnnnnn*—F—►

dbd_name

The DBD name as specified during DBDGEN.

nnnnnnnn

The number you specify of expected database segments for this MSDB. This number must be equal to or great than the number of MSDB segments loaded during restart.

The NBRSEGS parameter is also used to reserve space for terminal-related dynamic MSDBs for which no data has to be initially loaded.

F The optional page-fix indicator for this MSDB.

If the MSDBs are so critical to your Fast Path applications that IMS should not run without them, place a first card image at the beginning of the DBFMSDBx member. For each card image, the characters “MSDBABND=*n*” must be typed without blanks, and all characters must be within columns 1 and 72 of the card image. Four possible card images exist, and each contains one of the following sets of characters:

MSDBABND=Y

This card image causes the IMS control region to abend if an error occurs while loading the MSDBs during system initialization. Errors include:

- Open failure on the MSDBINIT data set
- Error in the MSDB definition
- I/O error on the MSDBINIT data set

MSDBABND=C

This card image causes the IMS control region to abend if an error occurs while writing the MSDBs to the MSDBCP1 or MSDBCP2 data set in the initial checkpoint after IMS startup.

MSDBABND=I

This card image causes the IMS control region to abend if an error occurs during the initial load of the MSDBs from the MSDBINIT data set, making one or more of the MSDBs unusable. These errors include data errors in the MSDBINIT data set, no segments in the MSDBINIT data set for a defined MSDB, and those errors described under “MSDBABND=Y.”

MSDBABND=A

This card image causes the IMS control region to abend if an error occurs during the writing of the MSDBs to the MSDBCPn data set (described in “MSDBABND=C”), or during the initial load of the MSDBs from the MSDBINIT data set (described in “MSDBABND=I”).

MSDBABND=B

This card image causes the IMS control region to abend if an error occurs during the writing of the MSDBs to the MSDBCPn data set (described in “MSDBABND=C”), or during the loading of the MSDBs in system initialization (described in “MSDBABND=Y”).

Auxiliary Storage Requirements for an MSDB

DASD space is needed to keep image copies of MSDBs when they are dumped at system and shutdown checkpoints. The data sets involved are the MSDBCP1 and MSDBCP2 data sets. The same calculations apply to the MSDBDUMP data set, which contains a copy of the MSDBs following a /DBDUMP DATABASE MSDB command.

The data sets just discussed are written in 2K-byte blocks. Because only the first extent is used, the allocation of space must be on cylinder boundaries and be contiguous.

Space allocation is calculated like this:

$SPACE = (2048, (R), , CONTIG, ROUND)$

The calculation of the number of records (R) to be allocated can be derived from the formula:

$$(E + P + 2047) / 2048$$

where:

E = main storage required, in bytes, for the Fast Path extension of the CNTs (ECNTs)

P = main storage required for all MSDBs as defined by the PROCLIB member DBFMSDBx

E is determined by the following formula:

$$E = (20 + 4D)T$$

where:

D = number of MSDBs using logical terminal names as keys

T = total number of logical terminal names defined in the system

High-Speed Sequential Processing (HSSP)

High-Speed Sequential Processing (HSSP) is a function of Fast Path that handles sequential processing of DEDBs.

Why HSSP?

Some reasons you may choose to use it are that, HSSP:

- Generally has a faster response time than regular batch processing.
- Optimizes sequential processing of DEDBs.
- Reduces program execution time.
- Typically produces less output than regular batch processing.
- Reduces DEDB updates and image copy operation times.
- Image copies can assist in database recovery.
- Locks at UOW level to ease “bottle-necking” of cross IRLM communication.
- Uses private buffer pools reducing impact on NBA/OBA buffers.
- Allows for execution in both a mixed mode environment, concurrently with other programs, and in an IRLM-using global sharing environment.
- Optimizes database maintenance by allowing the use of the image-copy option for an updated database.

More detailed information is included in the following topics on HSSP:

- “Limitations and Restrictions When Using HSSP”
- “Using HSSP” on page 285
- “HSSP Processing Option H (PROCOPT=H)” on page 285

Limitations and Restrictions When Using HSSP

Though HSSP can execute in a mixed-mode environment as well as concurrently with other programs, and in an environment with global sharing using IRLM; a program using HSSP can only execute as a non-message-driven BMP.

Other restrictions and limitations of HSSP include:

- Only one HSSP process can be active on an area at any given time. The /DIS AREA command identifies the IMSID of any HSSP job processing an area.
- HSSP processing and online utilities cannot process on the same area concurrently.
- Non-forward referencing while using HSSP is not allowed.
- Programs using HSSP must properly process the 'GC' status code by following it with a commit process.

Restrictions and limitations involving image copies include:

- The image copy option is available only for HSSP processing.
- HSSP image copying is allowed only if PROCOPT = H.
- The image copy process can only be done if a database is registered with DBRC. In addition, image copy data sets must be initialized in DBRC.

The following restrictions and limitations apply for PROCOPT=H:

- PROCOPT=H is allowed only for DEDBs.
- PROCOPT=H is not allowed on the segment level, only on the PCB level.
- Backward referencing while using HSSP is not allowed. You cannot use an HSSP PCB to refer to a prior UOW in a DEDB.
- Only one PROCOPT=H PCB per database per PSB is allowed.
- A maximum of four PROCOPTs can be specified, including H.

- PROCOPT=H must be used with other Fast Path processing options, such as GH and IH.
- When a GC status code is returned, the program must cause a commit process before any other call can be made to that PCB.
- HSSP image copying is not allowed if PROCOPT \neq H.
- An ACBGEN must be done to activate the PROCOPT=H.
- H is compatible with all other PROCOPTs except for PROCOPT=O.

Using HSSP

To use HSSP, you must specify a new PROCOPT option during PSBGEN, option 'H' see “HSSP Processing Option H (PROCOPT=H).” Additionally, you need to make sure that the programs using HSSP properly process the 'GC' status code by following it with a commit process.

HSSP includes the image-copy option and the ability to set area ranges. To use these functions, you need one or more of the following:

- The SETR statement
- The SETO statement
- A DFSCTL data set for the dependent regions
- DBRC
- PROCOPT=H

Related Reading: For more information about the SETR and SETO control statements, refer to *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

HSSP Processing Option H (PROCOPT=H)

PROCOPT=H is a PSBGEN OPTION. It allows you to define whether processing, with respect to a PCB, should be treated as an HSSP process. Its use provides HSSP capability for the application program using this PSB. Following is an example of macros and keywords for a PSBGEN using PROCOPT=H:

```
Label      PCB TYPE = DB
           ,DBDNAME = name
           ,PROCOPT = AH
```

Label is an optional parameter of the PCB macro. It can be up to 8 characters long and is identical to the label on the associated SETO or SETR statements. H is compatible with any other Fast Path PROCOPT, except for PROCOPT=O, and PROCOPT=H can be used in one or more PCBs.

Related Reading:

- For information on PROCOPT=H rules, see “Limitations and Restrictions When Using HSSP” on page 284.
- For more information on H processing, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Image-Copy Option

Selecting the image-copy option with HSSP reduces the total elapsed times of DEDB updates and subsequent image-copy operations.

High-Speed Sequential Processing (HSSP)

As database administrator, you decide whether to make an image copy of a database using HSSP. If you specify image copying, HSSP creates an asynchronous copy that is similar to a concurrent image copy.

The image copy process can only be done if a database is registered with DBRC. In addition, image copy data sets must be initialized in DBRC.

HSSP image copies can also be used for database recovery. However, the Database Recovery Utility must know that an HSSP image copy is supplied.

Related Reading: For information on DBRC databases and HSSP, and on created image copies, refer to the *IMS Version 9: Operations Guide* and the *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

For information on image copies and recovery, refer to *IMS Version 9: Utilities Reference: System*.

UOW Locking

In a globally shared environment, data is shared not only between IMS subsystems, but also across central processor complexes (CPC). In such an environment, communication between two IRLMs could potentially “bottleneck” and become impeded. To ease this problem, HSSP locks at a UOW level in update mode, reducing the locking overhead. Non-HSSP or DEDB online processing locks at a UOW level in a shared mode. Otherwise, the locking for DEDB online processing is at the CI level. For information on UOW locking, refer to *IMS Version 9: Administration Guide: System*.

Private Buffer Pools

Private buffer pools for the HSSP area are used for HSSP updates and image copies. HSSP does *not* impact NBA/OBA buffers. HSSP dynamically allocates up to three times the number of CIs per area in one UOW. Each buffer is a CI in size. The private buffer pools are located in ECSA/CSA.

HSSP jobs use a combination of both Private buffer pools and common buffers (NBA/OBA). HSSP dynamically allocates up to three times the number of CIs per area in one UOW, with each buffer being a CI in size. The private buffer pools are located in ECSA/CSA. HSSP uses the private buffers for reading RAP CIs, and common buffers for reading IOVF CIs. An FW status code may be received during the run of an HSSP job when NBA has been exceeded just as in a non-HSSP job.

Designing a DEDB or MSDB Buffer Pool

Buffers needed to fulfill requests resulting from database calls are obtained from a global pool called the Fast Path buffer pool. The characteristics of the pool are defined at IMS definition time and can be overridden at IMS startup time.

Three parameters characterize the Fast Path buffer pool:

DBBF

Total number of buffers.

The buffer pool is allocated at IMS startup in the ECSA or, if FPBUFF=LOCAL is specified in DFSFDRxx, in the FDBR private region. During emergency restart processing, the entire buffer pool can be briefly page-fixed. Consider the amount of available real storage when setting the DBBF value. IMS writes the total number of buffers to the X'5937' log.

For more information about page-fixing IMS resources, see *IMS Version 9: Administration Guide: System*.

DBFX

System buffer allocation.

This is a set of buffers in the Fast Path buffer pool that is page fixed at startup of the first region with access to Fast Path resources.

BSIZ

Buffer size.

The size must be larger than or equal to the size of the largest CI of any DEDB to be processed. The buffer size can be up to 28 KB.

Buffer Requirements

Fast Path buffers are used to hold:

- Update information such as:
 - MSDB FLD/VERIFY call logic
 - MSDB FLD/CHANGE call logic
 - MSDB updates (results of REPL, ISRT, and DLET calls)
 - Inserted SDEP segments
- Referenced DEDB CIs from the root addressable part and the sequential dependent part.
- Updated DEDB CIs from the root addressable part.
- SDEP segments that have gone through sync point. The SDEP segments are collected in the current SDEP segment buffer. One such buffer allocated for each area defined with the SDEP segment type exists. This allocation takes place at area open time.

The number of buffers a transaction or a sync interval is allowed to use must be specified for each region if Fast Path resources are likely to be accessed.

Normal Buffer Allocation (NBA)

Fast Path regions and IMS regions accessing Fast Path resources require that the normal buffer allocation (NBA) be specified in the region startup procedure. Because this allocation of buffers is used first, calculate them to accommodate most of the transaction requirements. At the start of the region, the number of NBA buffers is page fixed in the Fast Path buffer pool.

Overflow Buffer Allocation (OBA)

The overflow buffer allocation (OBA) is optional and is used for exceptional buffer requirements when the normal buffer allocation (NBA) has been exhausted. Its use is dependent on obtaining a latch that serializes all regions currently in an overflow buffer state. If the latch is not available, the region has to wait until it is available. After the latch has been obtained, the NBA value is increased by the OBA value and normal processing resumes. The overflow buffer latch is released during sync point processing. At any point in time, only the largest OBA request among all the active regions is page fixed in the Fast Path buffer pool.

Fast Path Buffer Allocation Algorithm

Fast Path buffers are allocated on demand up to a limit specified at the start of the region. Buffers so specified are called NBA to be used by one sync point interval.

Designing a DEDB or MSDB Buffer Pool

Before satisfying any request from the NBA allocation, an attempt is made to reuse any already allocated buffer containing an SDEP CI. This process goes on until the NBA limit is reached. From that point on, a warning in the form of an 'FW' status code returned to Fast Path database calls is sent to BMP regions. MD and MPP regions do not get this warning.

The next request for an additional buffer causes the buffer stealing facility to be invoked and then the algorithm examines each buffer and CI already allocated. As a result, buffers containing CIs being released are sent to a local queue (SDEP buffer chain) to be reused by this sync interval.

If, after invoking the buffer stealing facility, no available buffer is found, a request for the overflow buffer latch is issued. The overflow buffer latch governs the use of an additional buffer allocation called overflow buffer allocation (OBA). This allocation is also specified as a parameter at region start time. From that point on, any time a request cannot be satisfied locally, a buffer is acquired from the OBA allocation until the OBA limit is reached. At that time, MD and BMP regions have their 'FW' status code replaced by an 'FR' status code after an internal ROLB call is performed. In MD and MPP regions, the transaction is abended and stopped.

System Buffer Allocation (DBFX)

The system buffer allocation (DBFX) is needed, because DEDB writes are deferred until after sync point processing. The result of one transaction or sync interval is written back by one output thread. These output threads run from the control region in SRB mode. Buffers allocated to an output thread are therefore not available to dependent regions until after the CI they contain is written back. If the Fast Path buffer pool is defined exactly as the sum of all NBAs, dependent regions must wait for the buffers to come back to the global pool. Fast Path regions can process the next transaction as soon as the sync point completes. Sync point processing does not wait for the output thread to complete. The DBFX allocation of buffers is page fixed at the start of the first region specifying an NBA request.

Determining the Fast Path Buffer Pool Size

The number of Fast Path buffers (DBBFs) required is calculated using the following formula:

$$\text{DBBF} \geq A + N + \text{OBA} + \text{DBFX}$$

where:

- DBBF: Fast Path buffer pool size as specified
- A: Number of active areas that have SDEP segments
- NBA: Normal buffer allocation of each active region
- N: Total of all NBAs
- OBA: Largest overflow buffer allocation
- DBFX: System buffer allocation

Fast Path Buffer Performance Considerations

An incorrect specification of DBBF (too small) can result in the rejection of an area open or a region initialization. The system calculates the size of the buffer pool in accordance with the formula given in "Determining the Fast Path Buffer Pool Size" and rejects the open or initialization if the actual DBBF value is smaller.

A DBFX value that is too small is likely to cause region waits and increase response time.

An NBA value that is too small might cause the region processing to be serialized through the overflow buffer latch and again cause delays.

An NBA value that is too large can increase the probability of contention (and delays) for other transactions. All CIs can be acquired at the exclusive level and be kept at that level until the buffer stealing facility is invoked. This occurrence happens after the NBA limit is reached. Therefore, an NBA that is too large can increase resource contention.

A (NBA + OBA) value that is too small might result in more frequent unsuccessful processing. This means an 'FR' status code condition for BMP regions, or transactionabend for MD and MPP regions.

Inquiry-only programs do not make use of an OBA specification, as buffers already allocated are reused when the NBA limit is reached.

IMS logs information about buffers and their use to the X'5937' log. This information can be helpful in determining how efficiently the Fast Path buffers are being used.

The NBA Limit and Sync Point

In BMP regions, when the NBA limit is reached, an 'FW' status code is returned. This status code is presented to every subsequent Fast Path database call until the OBA limit condition is reached.

The first occurrence of the 'FW' status code indicates no more NBA buffers exist. This occurrence is a convenient point at which to request a sync point. Fast Path resources (and others) would be released and the next sync point interval would be authorized to use a new set of NBA buffers. The overflow buffer latch serializes all the regions in an overflow buffer state and therefore causes delays in their processing.

If processing is primarily sequential, the sync point should be invoked on a UOW boundary crossing.

Related Reading: See "Processing Option P (PROCOPT=P)" on page 275 for details on what happens on a UOW boundary crossing.

The DBFX Value and the Low Activity Environment

If the IMS or Fast Path activity in the system is relatively low, log buffers are written less often, and therefore output threads are scheduled or dispatched less frequently. This situation is likely to result in many buffers waiting to be written and therefore could cause wait-for-buffer conditions. Wait-for-buffer conditions could be alleviated by specifying a larger DBFX value.

A special case to be considered is the BMP region loading or processing a DEDB and being the only activity in the system. For example, assume an NBA of 20 buffers exists. To avoid a wait-for-buffer condition, the DBFX value must be specified as between one or two times the NBA value. This can result in a DBBF specification of three times the NBA number, which gives 60 buffers to the Fast Path buffer pool.

Except for the following case, there is no buffer look-aside capability across transactions or sync intervals (global buffer look-aside).

Designing a DEDB or MSDB Buffer Pool

Assume that a region requests a DEDB CI resource that is currently being written or is owned by another region that ends up being written (output thread processing). Then, this CI and the buffer are passed to the requestor after the write (no read required) completes successfully. Any other regions must read it from disk.

Designing a DEDB Buffer Pool in the DBCTL Environment

Buffers needed to fulfill requests from database calls are obtained from a global pool called the Fast Path buffer pool. The characteristics of the pool are defined at IMS definition time and can be overridden at IMS startup time.

Three parameters characterize the Fast Path buffer pool:

1. DBBF: Total number of buffers.

The buffer pool is allocated at IMS startup in the ECSA or, if FPBUFF=LOCAL is specified in DFSFDRxx, in the FDBR private region. IMS writes the total number of buffers to the X'5937' log.

2. DBFX: System buffer allocation.

This is a set of buffers in the Fast Path buffer pool that is page fixed at startup of the first region with access to Fast Path resources.

3. BSIZ: Buffer size.

The size must be larger than or equal to the size of the largest CI of any DEDB to be processed. The buffer size can be up to 28 KB.

Buffer Requirements in a DBCTL Environment

Fast Path buffers are used to hold:

- Update information such as inserted SDEP segments.
- Referenced DEDB CIs from the root addressable part and the sequential dependent part.
- Updated DEDB CIs from the root addressable part.
- SDEP segments that have gone through sync point. The segments are collected in the current SDEP segment buffer. One buffer allocated for each area defined with the SDEP segment type exists. This allocation takes place at area open time.

The number of buffers a transaction or a sync interval is allowed to use must be specified for each region if Fast Path resources are likely to be accessed.

Normal Buffer Allocation for BMPs

BMP regions accessing Fast Path resources require this allocation to be specified in the region startup procedure. The startup parameter is already specified as NBA. This allocation of buffers is used first and should be calculated to accommodate most of the transaction requirements. At the start of the region, the number of NBA buffers is page fixed in the Fast Path buffer pool.

Normal Buffer Allocation for CCTL Regions and Threads

CCTL (coordinator control) regions, requiring Fast Path resources, need the following parameters specified in the DRA startup table:

- CNBA
- FPB

CNBA is the normal buffer allocation of each active CCTL region. FPB is the normal buffer allocation for CCTL threads.

When the CCTL connects to DBCTL, the number of CNBA buffers is page fixed in the Fast Path buffer pool. However, if CNBA buffers are not available, the connect fails.

Each CCTL thread that requires DEDB buffers is assigned its Fast Path buffers (FPB) out of the total number of CNBA buffers.

For more information about the CCTLNBA parameter, refer to *IMS Version 9: Administration Guide: System*.

Overflow Buffer Allocation for BMPs

This buffer allocation is optional and is used for exceptional buffer requirements when the NBA has been exhausted. Its use is dependent on obtaining a latch that serializes all BMPs and CCTL threads currently in an overflow buffer state. If the latch is not available, the region has to wait until it is available. After the latch has been obtained, the NBA value is increased by the OBA value and normal processing resumes. The overflow buffer latch is released during sync point processing. At any point in time, only the largest OBA request among all the active BMPs and CCTL threads is page fixed in the Fast Path buffer pool.

Overflow Buffer Allocation for CCTL Threads

OBA for CCTL threads is similar to that for BMPs. The OBA value used for each thread is set with the FPOB parameter in the startup table. This buffer allocation is optional and is used for exceptional buffer requirements when the FPB has been exhausted. Its use is dependent on obtaining a latch that serializes all BMPs and CCTL threads currently in an overflow buffer state. If the latch is not obtained, the FPB value is increased by the FPOB value, and normal processing resumes. The overflow buffer latch is released during sync point processing. At any point in time, only the largest OBA/FPOB request among all the active BMPs and CCTL threads is page fixed in the Fast Path buffer pool.

Fast Path Buffer Allocation Algorithm for BMPs

FPBs are allocated on demand up to a limit specified at the start of the region. Buffers specified as NBAs are used by one sync point interval.

Before satisfying any request from the NBA allocation, an attempt is made to reuse any already allocated buffer containing an SDEP CI. This process goes on until the NBA limit is reached. From that point on, a warning in the form of an 'FW' status code returned to Fast Path database calls is sent to BMP regions.

The next request for an additional buffer causes the buffer stealing facility to be invoked and then the algorithm examines each buffer and CI already allocated. As a result, buffers containing CIs being released are sent to a local queue (SDEP buffer chain) to be reused by this sync interval.

If, after invoking the buffer stealing facility, no available buffer is found, a request for the overflow buffer latch is issued. The overflow buffer latch governs the use of an additional buffer allocation, OBA. This allocation is also specified as a parameter at region start time. From that point on, any time a request cannot be satisfied locally,

Designing a DEDB Buffer Pool in the DBCTL Environment

a buffer is acquired from the OBA allocation until the OBA limit is reached. At that time, BMP regions have their 'FW' status code replaced by an 'FR' status code after an internal ROLB call is performed.

Fast Path Buffer Allocation Algorithm for CCTL Threads

When a CCTL thread issues a schedule request using FPB, buffers are allocated out of the CNBA total. If FPB cannot be satisfied out of CNBA, the schedule request fails.

Before satisfying any request from the FPB allocation, an attempt is made to reuse any already allocated buffer containing an SDEP CI. This process goes on until the FPB limit is reached. From that point on, a warning in the form of an 'FW' status code returned to Fast Path database calls is sent to the CCTL threads.

The next request for an additional buffer causes the buffer stealing facility to be invoked, and then the algorithm examines each buffer and CI already allocated. As a result, buffers containing CIs being released are sent to a local queue (SDEP buffer chain) to be reused by this sync interval.

If, after invoking the buffer stealing facility, no available buffer is found, a request for the overflow buffer latch is issued. The overflow buffer latch governs the use of an additional buffer allocation, OBA (FPOB). From that point on, any time a request cannot be satisfied locally, a buffer is acquired from the FPOB allocation until the FPOB limit is reached. At that time, CCTL threads have their 'FW' status code replaced by an 'FR' status code after an internal ROLB call is performed.

System Buffer Allocation (SBA)

The system buffer allocation (SBA) is needed because DEDB writes are deferred until after sync point processing. The result of one sync interval is written back by one output thread. These output threads run from the control region in SRB mode. Buffers allocated to an output thread are therefore not available to BMPs and CCTL threads until after the CI they contain is written back. If the Fast Path buffer pool is defined exactly as the sum of all NBAs, BMPs and CCTL threads must wait for the buffers to come back to the global pool. BMPs and CCTL threads can process the next transaction as soon as the sync point completes. Sync point processing does not wait for the output thread to complete. The DBFX allocation of buffers is page fixed at the start of the first region specifying an NBA or FPB request.

Determining the Size of the Fast Path Buffer Pool for DBCTL

The number of buffers required is calculated using the following formula:

$$DBBF \geq A + N + LO + DBFX + CN$$

Where the values are:

- DBBF: Fast Path buffer pool size as specified
- A: Number of active areas that have SDEP segments
- N: Total of all NBAs
- LO: Largest overflow buffer allocation among active BMPs and CCTL threads
- DBFX: System buffer allocation
- CN: Total of all CNBAs

Fast Path Buffer Performance Considerations for DBCTL

An incorrect specification of DBBF (too small) can result in the rejection of an area open or a region initialization. The system calculates the size of the buffer pool in accordance with the formula given in “Determining the Size of the Fast Path Buffer Pool for DBCTL” on page 292 and rejects the open or initialization if the actual DBBF value is smaller.

A DBFX value that is too small is likely to cause region waits and increase response time.

An NBA/FPB value that is too small might cause the region processing to be serialized through the overflow buffer latch and again cause delays.

An NBA/FPB value that is too large can increase the probability of contention (and delays) for other BMPs and CCTL threads. All CIs can be acquired at the exclusive level and be kept at that level until the buffer stealing facility is invoked. This happens after the NBA limit is reached. Therefore, an NBA/FPB that is too large can increase resource contention. Also, an FPB value that is too large indicates that fewer CCTL threads can concurrently schedule Fast Path PSBs.

A (NBA + OBA) value that is too small might result in more frequent unsuccessful processing. This means an 'FR' status code condition for BMP regions and CCTL threads.

Inquiry-only BMP or CCTL programs do not make use of the overflow buffer specification logic, as buffers already allocated are reused when the NBA/FPB limit is reached.

IMS logs information about buffers and their use to the X'5937' log. This information can be helpful in determining how efficiently the Fast Path buffers are being used.

The NBA/FPB Limit and Sync Point in a DBCTL Environment

In BMP regions and CCTL threads, when the NBA/FPB limit is reached, an 'FW' status code is returned. This status code is presented to every subsequent Fast Path database call until the OBA/FPOB limit condition is reached.

The first occurrence of the 'FW' status code indicates no more NBA/FPB buffers exist. This occurrence is a convenient point at which to request a sync point. Fast Path resources (and others) would be released and the next sync point interval would be authorized to use a new set of NBA/FPB buffers. The overflow buffer latch serializes all the regions in an overflow buffer state and therefore causes delays in their processing.

Related Reading: See “Processing Option P (PROCOPT=P)” on page 275 for benefits of using PROCOPT=P for BMP regions.

Low Activity and the DBFX Value in a DBCTL Environment

If the IMS or Fast Path activity in the system is relatively low, log buffers are written less often and therefore output threads are scheduled or dispatched less frequently. This situation is likely to result in many buffers waiting to be written and therefore could cause wait-for-buffer conditions. This could be alleviated by specifying a larger DBFX value.

Designing a DEDB Buffer Pool in the DBCTL Environment

Consider the special case: The BMP region loads or processes a DEDB and is the only activity in the system. For example, assume that an NBA of 20 buffers exists. To avoid a wait-for-buffer condition, the DBFX value must be between once or twice the NBA value. This can result in a DBBF specification of three times the NBA number, giving 60 buffers to the Fast Path buffer pool.

Except for the following case, there is no buffer look-aside capability across BMP regions and CCTL threads or sync intervals (global buffer look-aside).

Assume that a region requests a DEDB CI resource that is currently being written or is owned by another region that ends up being written (output thread processing). Then, this CI and the buffer are passed to the requestor after the successful completion of the write (no read required). Any other BMP regions and CCTL threads must read it from disk.

A Note on Fast Path Buffer Allocation in IMS Regions

IMS regions that access Fast Path resources must have the NBA and OBA parameters specified in their startup procedures.

With MODE=MULT, these allocations must be large enough to accommodate all buffer requirements for transactions processed between sync points.

With MODE=SNGL, transaction classes should be set up so transactions with similar buffer requirements are run in the same region.

Chapter 11. Implementing Database Design

After you have designed your databases and before application programs can use them, you must tell IMS their physical and logical characteristics by coding and generating a DBD (database description) for each database.

Before an application program can use the database, you must tell IMS the application program's characteristics and use of data and terminals. You tell IMS the application program characteristics by coding and generating a PSB (program specification block).

Finally, before an application program can be scheduled for execution, IMS needs the PSB and DBD information for the application program available in a special internal format called an ACB (application control block).

This chapter examines the following areas of implementing your database design:

- "Coding Database Descriptions as Input for the DBDGEN Utility"
- "Implementing HALDB Design" on page 298
- "Coding Program Specification Blocks as Input to the PSBGEN Utility" on page 305
- "Building the Application Control Blocks (ACBGEN)" on page 308
- "Defining Generated Program Specification Blocks for SQL Applications" on page 309

Coding Database Descriptions as Input for the DBDGEN Utility

A DBD is a series of macro instructions that describes such things as a database's organization and access method, the segments and fields in a database record, and the relationships between types of segments. After you have coded the DBD macro instructions, they are used as input to the DBDGEN utility. This utility is a macro assembler that generates a DBD control block and stores it in the IMS.DBDLIB library for subsequent use during database processing.

Figure 174 illustrates the DBD generation process. Figure 175 on page 296 shows the input to the DBDGEN utility. Separate input is required for each database being defined.

Coding DBDs for DBDGEN Utility

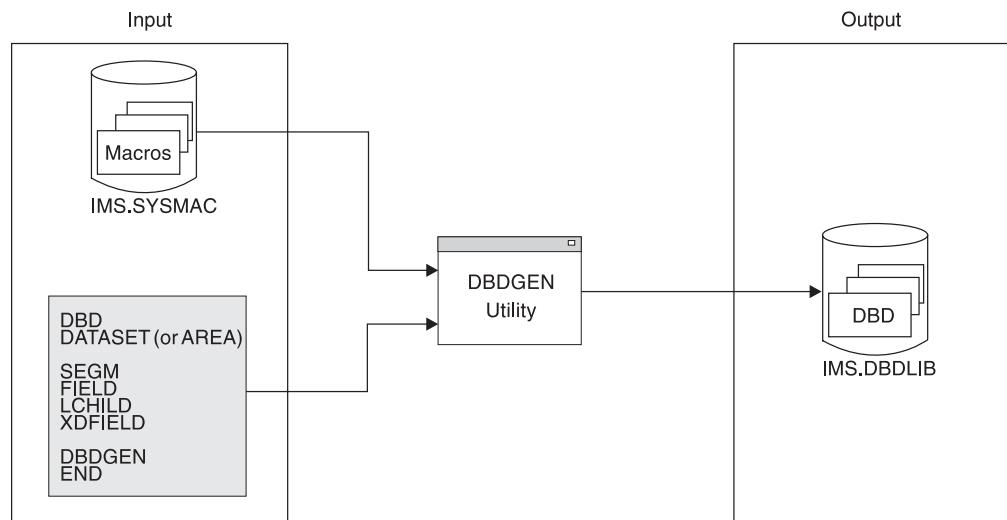


Figure 174. The DBD Generation Process

```
//DBDGEN      JOB  MSGLEVEL=1
//              EXEC      DBDGEN,MBR=APPLPGM1
//C.SYSIN      DD      *
```

| | |
|-------------------|---|
| DBD | required for each DBD generation |
| data set(or AREA) | required for each data set group (or AREA in a Fast Path DEDB) |
| SEGM | required for each segment type |
| FIELD | required for each DBD generation |
| LCHILD | required for each secondary index or logical relationship |
| XDFIELD | required for each secondary index relationship |
| . | |
| . | |
| . | |
| DBDGEN | required for each DBD generation |
| END | required for each DBD generation |

```
/*
```

Figure 175. Structure of DBD Generation Input

The DBD Statement

In the input, the DBD statement names the database being described and specifies its organization. Only one DBD statement exists in the input deck.

The DATASET Statement

This statement defines the physical characteristics of the data sets to be used for the database. At least one DATASET statement is required for each data set group in the database. Depending on the type of database, up to 10 data set groups can be defined. Each DATASET statement is followed by the SEGM statements for all segments to be placed in that data set group.

The DATASET statement is not allowed for HALDBs. Use either the HALDB Partition Definition utility to define HALDB partitions or the DBRC commands INIT.DB and INIT.PART

If the database is a DEDB, the AREA statement is used instead of the DATASET statement. The AREA statement defines an area in the DEDB. Up to 2048 AREA

statements can be used to define multiple areas in the database. All AREA statements must be put between the DBD statement and the first SEGM statement.

The SEGM Statement

This statement defines a segment type in the database, its position in the hierarchy, its physical characteristics, and its relationship to other segments. SEGM statements are put in the input deck in hierarchic sequence, and a maximum of 15 hierarchic levels can be defined. The number of database statements allowed depends on the type of database. SEGM statements must immediately follow the data set or AREA statements to which they are related.

The FIELD Statement

This statement defines a field within a segment type. FIELD statements must immediately follow the SEGM statement to which they are related. A FIELD statement is required for all sequence fields in a segment and all fields the application program can refer to in the SSA of a DL/I call. A FIELD statement is also required for any fields referenced by a SENFLD statement in any PSB. To save space, do not generate FIELD statements except in these circumstances. FIELD statements can be put in the input deck in any order except that the sequence field, if one is defined, must always be first. Up to 255 fields can be defined for each segment type, and a maximum of 1000 fields can be defined for each database. The definition of fields within a segment can overlap. For example, a date “field” within a segment can be defined as three 2-byte fields and also as one 6-byte field as shown in Figure 176.

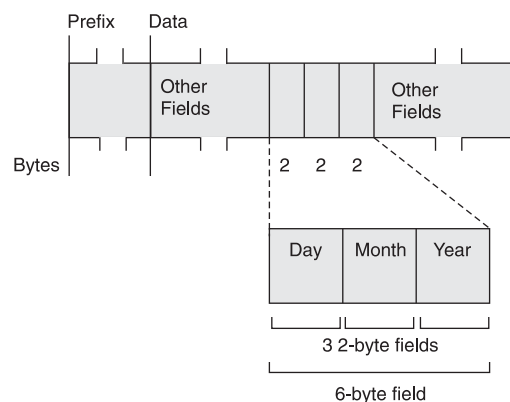


Figure 176. Example of a Date Field within a Segment Defined as Three 2-Byte Fields and One 6-Byte Field

This technique allows application programs to access the same piece of data in a variety of ways. To access the same piece of data in a variety of ways, you code a separate FIELD statement for each field. For the example shown, you would code four FIELD statements, one for the total 6-byte date and three for each 2-byte field in the date.

The LCHILD Statement

The LCHILD statement defines a secondary index or logical relationship between two segment types, or the relationship between a HIDAM (or PHIDAM) index database and the root segment type in the HIDAM (or PHIDAM) database. LCHILD statements immediately follow the SEGM, FIELD, or XDFLD statement of the segment involved in the relationship. Up to 255 LCHILD statements can be defined for each database.

Restriction: The LCHILD statement cannot be specified for the primary index of a PHIDAM database because the primary index is automatically generated.

The XDFLD Statement

The XDFLD statement is used only when a secondary index exists. It is associated with the target segment and specifies:

- The name of an indexed field
- The name of the source segment
- The field used to create the secondary index from the source segment

Up to 32 XDFLD statements can be defined per segment. However, the number of XDFLD and FIELD statements combined cannot exceed 255 per segment or 1000 per database.

Restriction: The CONST parameter is not allowed for a HALDB. Shared secondary indexes are not supported.

The DBDGEN and END Statements

One DBDGEN statement and one END statement is put at the end of each DBD generation input deck. These specify:

- The end of the statements used to define the DBD (DBDGEN)
- The end of input statements to the assembler (END)

Related Reading: Detailed instructions for coding DBD statements and examples of DBDs are contained in *IMS Version 9: Utilities Reference: System*.

Implementing HALDB Design

To define a HALDB, you define it to DBRC. You can do this using either DBRC batch commands or using the Partition Definition utility. This topic discusses defining HALDBs using the Partition Definition utility. This topic also discusses “Allocating an ILDS” on page 304.

Related Reading: *The Complete IMS HALDB Guide*, published by IBM Redbooks™ for the Version 8 release of IMS, contains a comprehensive discussion of HALDBs.

Creating HALDBs with the HALDB Partition Definition Utility

The HALDB Partition Definition utility is an ISPF application that allows you to manage IMS HALDB partitions.

The HALDB Partition Definition utility is accessed through ISPF panels in a TSO session. You can perform the following tasks on the HALDB master and its partitions:

- Register a new HALDB master database with DBRC.
- Add HALDB partitions to an existing HALDB.
- Find, view, sort, copy, modify, delete, and print HALDB partitions.
- Define and modify data set groups.
- Edit HALDB information.
- Export HALDB definitions.
- Import HALDB definitions.
- View IMS DDNAME concatenations.

- Choose IMS RECON and DBDLIB libraries.
- Delete HALDB information.

Creating HALDB Partitions With the Partition Definition Utility

To create a HALDB, you first use the Database Description Generation (DBDGEN) utility to create a master database and define it as a HALDB. After you have defined a HALDB master database, use the Partition Definition utility to define the partitions within the HALDB.

Related Reading: For information on using the DBDGEN utility to create a HALDB master database, see:

- Figure 161 on page 240 for an example of the DBD for PHDAM
- “Coding Database Descriptions as Input for the DBDGEN Utility” on page 295
- *IMS Version 9: Utilities Reference: System*

When you define the first HALDB partition, you must also register the HALDB master database in the DBRC RECON data set. You can use either the HALDB Partition Definition utility or the DBRC INIT.DB and INIT.PART commands to do this. The HALDB Partition Definition utility does not impact RECON data set contention of online IMS subsystems. The RECON data set is reserved only for the time it takes to process a DBRC request. It is not held for the duration of the utility execution.

Related Reading: For additional information on HALDB and the RECON data set, see *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

When defining HALDB partitions using the Partition Definition utility, you must provide information such as the partition name, data set prefix name, and high key value. Whenever possible, the Partition Definition utility provides default values for required fields.

The steps for defining a new HALDB are as follows:

1. Use the DBDGEN process to define a HALDB master database. The HALDB Partition Definition utility does not let you define HALDB partitions until the DBDGEN process is performed.
2. Make the dialog data sets available to the TSO user. You can add the data sets to a LOGON procedure or use TSO commands to allocate them. You can use the TSOLIB command to add data sets to the STEPLIB. Table 21 shows which file names and data sets need to be allocated. Be sure to use your own high level qualifiers.

Table 21. File Names and Data Sets to Allocate

| File Name | Sample Data Set Names | Disposition |
|-----------|-----------------------|-------------|
| STEPLIB | IVPEXE91.SDFSRESL | N/A |
| SYSPROC | IVPEXE91.SDFSEXEC | SHR |
| ISPLMLIB | IVPEXE91.SDFSMLIB | SHR |
| ISPLLIB | IVPEXE91.SDFSPLIB | SHR |
| ISPTLIB | IVPEXE91.SDFSTLIB | SHR |
| IMS | IVPEXE91.DBDLIB | SHR |

If you use a logon procedure, you must log on again and specify logon with the new procedure. If you use allocation commands, they must be issued outside of

HALDB Partition Utility

ISPF. After you allocate the data sets and restart ISPF, restart the Install/IVP dialog, return to this task description, and continue with the remaining steps.

3. Start the HALDB Partition Definition utility from the ISPF command line by issuing the following command:

```
TSO %DFSHALDB
```

You can use the F2 key to split the screen and view these instructions online while viewing the HALDB partition definition panels at the same time.

4. Specify the name of the database. Fill in the first partition name as shown in Figure 177. Fill in the data set name prefix using the data set name for your data set instead of the high level qualifier shown in Figure 177. You should, however, specify the last qualifier as IVPDB1A to match cluster names previously allocated.

Recommendation: When naming your partitions, use a naming sequence that allows you to add new names later without disrupting the sequence. For example, if you name your partitions xxxx010, xxxx020 and xxxx030 and then later split partition xxxx020 because it has grown too large, you can name the new partition xxxx025 without disrupting the order of your naming sequence.

```
Help
-----
                        Partition Default Information

Type the field values.  Then press Enter to continue.

Database Name . . . . . IVPDB1

                        Processing Options
Automatic Definition . . . .No
Input data set . . . . .
Use defaults for DS groups .No

                        Defaults for Partitions
Partition Name . . . . .IVPDB11
Data set name prefix . . . .IXUEXEHQ.IVPDB1A

Free Space
Free block freq. factor . 0
Free space percentage . . 0

                        Defaults for data set groups
Block Size . . . . .8192

DBRC options
Max. image copies . . . .2
Recovery period . . . .0
Recovery utility JCL . . RECOVJCL
Default JCL . . . . .
Image Copy JCL . . . . .ICJCL
Online image copy JCL . .0ICJCL
Receive JCL . . . . .RECVJCL
Reusable? . . . . .No

To exit the application, press F3

Command = = = >
```

Figure 177. Partition Default Information

5. Define your partitions in the Change Partition panel. Make sure that the name of the partition and the data set name prefix are correct and then define a high key for the partition.

The high key identifies the highest root key of any record that the partition can contain and is represented by a hexadecimal value that you enter directly into the Partition High Key field of the Change Partition panel. Press F5 to accept the hexadecimal value and display its alphanumeric equivalent in the right section of the Partition High Key field.

You can enter the partition high key value using alphanumeric characters by pressing F5 before making any changes in the hexadecimal section of the Partition High Key field. This displays the ISPF editing panel. The alphanumeric input you enter in the editing panel displays in both hexadecimal and alphanumeric formats in the Change Partition Panel when you press F3 to save and exit the ISPF editor.

The last partition you define for a HALDB must have a high key value of X'FF'. This ensures that the keys of all records entered into the HALDB will be lower than the highest high key in the HALDB. The Partition Definition utility fills all remaining bytes in the Partition High Key field with hexadecimal X'FF'.

When you finish defining the partition high key, press enter to create the partition. The Change Partition panel remains active so that you can create additional partitions. To create additional partitions, you must change the partition name and the partition high key.

Figure 178 on page 302. is an example of the Change Partition panel. The Partition High Key field includes sample input.

6. When you finish defining partitions, press the cancel key (F12) to exit the Change Partition panel. A list of partitions defined in the current session displays.

To exit the HALDB Partition Definition utility entirely, press F12 again.

```

Help
-----
                                Change Partition
Type the field values. Then press Enter.

Database name.....IVPDB1
Partition name.....IVPDB11
Partition ID.....1
Data set name prefix...IXUEXEHQ.IVPDB1A
Partition Status.....

Partition High Key
+00 57801850 00F7F4F2 40C5A585 99879985 | ...&.742 Evergre |
+10 859540E3 85999981 | en Terra |

Free Space
Free block freq. factor...0
Free space percentage.....0

Attributes for data set group A
Block Size.....8192

DBRC options
Max. image copies.....2
Recovery period.....0
Recovery utility JCL.....
Image copy JCL.....ICJCL
Online image copy JCL....OICJCL
Receive JCL.....RECVJCL
Reusable?.....No

Command  = = = >

```

Figure 178. Change Partition Panel

Automatic and Manual HALDB Partition Definition: You can choose either automatic or manual partition definition by specifying Yes or No in the Automatic Definition field in the Processing Options section of the Partition Default Information panel.

Entering Yes in the Automatic Definition field specifies that the Partition Definition utility automatically defines your HALDB partitions. You must have previously created a data set and it must contain your HALDB partition selection strings. Specify the name of the data set in the Input data set field.

Entering No in the Automatic Definition field specifies that you define your HALDB partitions manually. “Creating HALDB Partitions With the Partition Definition Utility” on page 299 explains this process. You can still use an input data set when you define HALDB partitions manually.

Adding HALDB Partitions to an Existing HALDB

Related Reading: See Appendix E, “HALDB Partition Definition utility,” on page 517 for information on using the Partition Definition utility for adding HALDB partitions to an existing HALDB.

Finding, Viewing, Sorting, Copying, Modifying, Deleting, and Printing HALDB Partitions

Related Reading: See Appendix E, “HALDB Partition Definition utility,” on page 517 for information on the interfaces for finding, viewing, sorting, copying, modifying, deleting, and printing HALDB partitions.

Defining and Modifying Data Set Groups

Related Reading: See Appendix E, “HALDB Partition Definition utility,” on page 517 for information on the interfaces for defining and modifying HALDB data set groups.

Number of Data Sets in a HALDB Partition

Table 22 lists the minimum and maximum number of data sets a partition can contain for each type of HALDB.

Table 22. Minimum and maximum number of data sets for HALDB partitions.

| HALDB Type | Minimum number of data sets | Maximum number of data sets |
|------------|--|--|
| PHDAM | Two: an OSAM or VSAM entry-sequenced data set (ESDS), and a key-sequenced data set (KSDS) for the ILDS | Eleven: ten OSAM or VSAM ESDS, and one KSDS for the ILDS |
| PHIDAM | Three: a OSAM or VSAM ESDS, a KSDS for the ILDS, and a KSDS for the primary index | Thirteen: ten OSAM or VSAM ESDS, one KSDS for the ILDS, and one KSDS for the primary index |
| PSINDEX | One: a KSDS | One: a KSDS |

Exporting Database Definitions

Related Reading: See Appendix E, “HALDB Partition Definition utility,” on page 517 for information on the interfaces for exporting Database definitions.

Use the HALDB Partition Definition utility to export a HALDB definition. The database information is stored in the partitioned data set that you specify as an ISPF table and so must have the attributes of ISPTLIB data sets (record format = fixed block, record length = 80, data set organization = PDS or PDS/E).

Importing Database Definitions

Only an exported ISPF table can be used for importing database definitions.

The output from the export of a HALDB is a member of a PDS. The information about the HALDB is saved in the form of an ISPF table. The ISPF table becomes input for the import process.

The import can be performed from the HALDB Partition Definition utility or a batch job.

To import a database using a batch job, submit a batch ISPF job similar to the job shown in Figure 337 on page 547. All ISPF DD names are required.

The batch job executes the standard ISPF command, ISPSTART, that sets up the ISPF environment, and then starts the DSPXRUN command. The DSPXRUN command identifies the database, the import file to use, and the processing options.

Related Reading: For more information on the DSPXRUN command, see “DSPXRUN Command Syntax” on page 548.

Viewing IMS DDNAME Concatenation

You can look at the concatenation of data sets that are allocated to the IMS DDNAME. The data set is displayed using the ISRDDN command, which is part of the ISPF product.

When you specify a generic database name and use options 1 through 5 from the DFSHALDB panel, the viewing IMS DDNAME concatenation option only works if

HALDB Partition Utility

you use 4 or fewer DBD data sets. If you specify option 7, the data sets concatenated to the IMS DDNAME always display.

Use the help (F1) information provided by ISRDDN and ISPF to learn more about the ISRDDN utility. When you exit the ISRDDN utility, you return to the HALDB Partition Definition utility panels.

Choosing IMS RECON and DBDLIB Libraries

The HALDB Partition Definition utility menu contains an option to set IMS configurations. The IMS configuration can consist of a combination of DBDLIB and RECON data sets. You should use the same DBDLIB and RECON data sets that IMS will use to access the database. You can specify one data set for RECON1, RECON2, RECON3, and up to ten DBDLIB data sets for the IMS DDNAME.

You can control the RECON data sets in a configuration. If you have the IMS DDNAME allocated from the logon procedure and the IMS.SDFSRESL libraries allocated to the STEPLIB DDNAME, do not use the configuration option. If you define and select a configuration, those data sets override the allocations from the logon procedure.

The IMS DDNAME includes the data sets that contain the DBDLIB members. The STEPLIB allocation contains the RECON1, RECON2, and RECON3 members that name the actual RECON data sets. The RECON/DBDLIB Configurations option re-allocates the IMS DDNAME and allocates RECON1, RECON2, and RECON3 DDNAMEs to specify the RECON data sets.

If you delete a configuration only, the configuration is deleted from the list, but the data sets that are named in the configuration are not deleted.

Deleting Database Information

Use the HALDB Partition Definition utility to delete a database. Use the / (slash) character to confirm that you want to delete the database. You may wish to perform an export prior to deleting a database from the RECON data set.

Attention: You cannot undo the delete of a HALDB database.

Allocating an ILDS

Partitioning a database can complicate the use of pointers between database records because after a partition has been reorganized the following pointers may become invalid:

- Pointers from other database records within this partition
- Pointers from other partitions that point to this partition
- Pointers from secondary indexes

The use of indirect pointers eliminates the need to update pointers throughout other database records when a single partition is reorganized. The Indirect List data set (ILDS) acts as a repository for the indirect pointers. There is one ILDS per partition in PHDAM and PHIDAM databases.

The ILDS contains indirect list entries (ILEs). Each ILE in an ILDS has a 9-byte key that is the indirect list key (ILK) of the target segment appended with the segment code of the target segment. The ILK is a unique token that is assigned to segments when the segments are created.

After a reorganization reload or a migration reload of segments involved in inter-record pointing, the ILE is updated to reflect the changes in location of the target segment of the ILE. Segments involved in inter-record pointing can be one of the following types:

- Physically paired logical children
- Logical parents of unidirectional logical children
- Targets of secondary indexes

The sample command in Figure 179 defines an ILDS. Note that the key size is 9 bytes at offset 0 (zero) into the logical record. Also note that the record size is specified as 50 bytes, the current length of an ILE.

```
DEFINE CLUSTER ( -
    NAME (FFDBPRT1.XABCD010.L00001) -
    TRK(2,1) -
    VOL(IMSQAV) -
    FREESPACE(80,10) -
    REUSE -
    SHAREOPTIONS(3,3) -
    SPEED ) -
DATA ( -
    NAME(FFDBPRT1.XABCD010.INDEXD) -
    CISZ(512) -
    KEYS(9,0) -
    RECSZ(50,50) ) -
INDEX ( -
    NAME(FFDBPRT1.XABCD010.INDEXS) -
    CISZ(2048) )
```

Figure 179. Sample Command to Define an ILDS

To compute the size of an ILDS, multiply the size of an ILE by the total number of physically paired logical children, logical parents of unidirectional relationships, and secondary index targets.

Related Reading:

- For information about the role of ILDS in the HALDB self-healing pointer process, see “The HALDB Self-Healing Pointer Process” on page 388.
- For information about initializing an ILDS, search for “Indirect List Data Set” in *IMS Version 9: Administration Guide: System*.

Coding Program Specification Blocks as Input to the PSBGEN Utility

A PSB is a series of macro instructions that describes an application program’s characteristics, its use of segments and fields within a database, and its use of logical terminals. A PSB consists of one or more PCBs (program communication blocks). Of the two types of PCBs, one is used for alternate message destinations, the other, for application access and operation definitions.

After you code the PSB macro instructions, they are used as input to the PSBGEN utility. This utility is a macro assembler that generates a PSB control block then stores it in the IMS.PSBLIB library for subsequent use during database processing.

Figure 180 on page 306 shows the PSB generation process.

Coding PSBs for the PSBGEN Utility

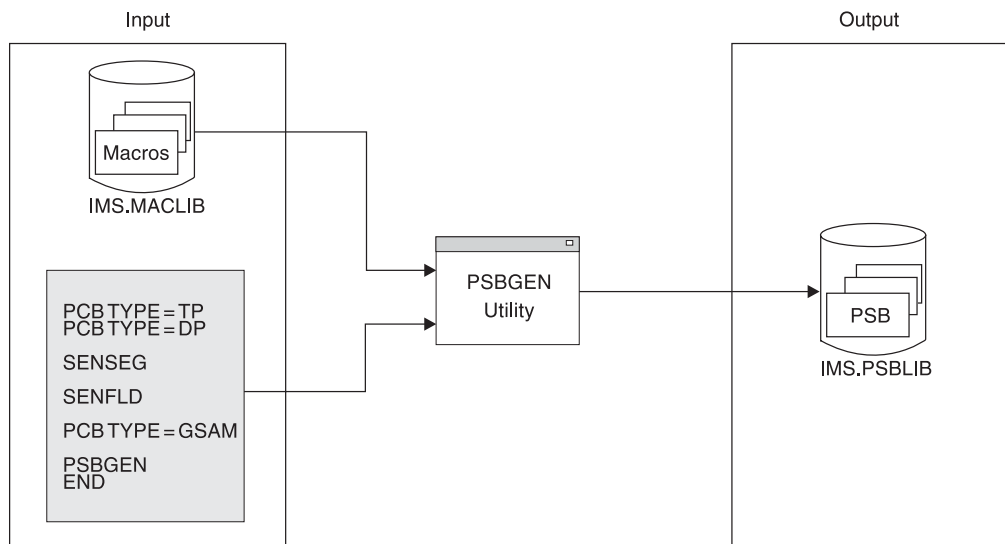


Figure 180. The PSB Generation Process

Figure 181 shows the structure of the deck used as input to the PSBGEN utility.

```
//PSBGEN    JOB  MSGLEVEL=1
//          EXEC      PSBGEN,MBR=APPLPGM1
//C.SYSIN    DD      *

      PCB TYPE=TP      required for output message destinations
      PCB TYPE=DB      required for each database the application program
                        can access
      SENSEG          required for each segment in the database the
                        application program can access
      SENFLD          required for each field in a segment that
                        the application program can access,
                        when field-level sensitivity is specified

      PCB TYPE=GSAM

      :
      :
      PSBGEN          required for each PSB generation
      END              required for each PSB generation
/*
```

Figure 181. Structure of PSB Generation Input

The Alternate PCB Statement

Two types of PCB statements can be placed in the input deck: the *alternate* PCB statement and the *database* PCB statement. The alternate PCB statement describes where a message can be sent when the message's destination differs from the place where it was entered. Alternate PCB statements must be put at the beginning of the input deck. More information on alternate PCBs is contained in *IMS Version 9: Administration Guide: System*.

The Database PCB Statement

The database PCB statement defines the DBD of the database that the application program will access. Database PCB statements also define types of operations (such as get, insert, and replace) that the application program can perform on segments in the database. The database can be either physical or logical. A separate database PCB statement is required for each database that the application program accesses. In each PSB generation, up to 2500 database PCBs can be

defined, minus the number of alternate PCBs defined in the input deck. The other forms of statements that apply to PSBs are SENSEG, SENFLD, PSBGEN, and END.

The SENSEG Statement

This statement defines a segment type in the database to which the application program is sensitive. A separate SENSEG statement must exist for each segment type. The segments can physically exist in one database or be derived from several physical databases. If an application program is sensitive to a segment beneath the root segment, it must also be sensitive to all segments in the path from the root segment to the sensitive segment. For example, in Figure 182 if D is defined as a sensitive segment for an application program, B and A must also be defined as sensitive segments.




Figure 182. Example of a SENSEG Relationship

An application program must be sensitive to all segments in the path to the segment that you actually want to be sensitive. However, you can make the application program sensitive to only the segment key in these other segments. With this option, the application program does not have any access to the segments other than the keys it needs to get to the sensitive segment. To make an application sensitive to only the segment key of a segment, code PROCOPT=K in the SENSEG statement. The application program will not be able to access any other field in the segment other than the segment key. In the previous example, the application program would be sensitive to the key of segment A and B but not sensitive to A and B's data.

SENSEG statements must immediately follow the PCB statement to which they are related. Up to 30000 SENSEG statements can be defined for each PSB generation.

The SENFLD Statement

This statement is used only in parallel with field-level sensitivity. It defines the fields in a segment type to which the application program is sensitive. This statement, in conjunction with the SENSEG statement, helps you secure your data. Each SENFLD statement must follow the SENSEG statement to which it is related. Up to 255 sensitive fields can be defined for a given segment type, and a maximum of 10000 can be defined for each PSB generation.

The PSBGEN Statement

This statement names the PSB and specifies various characteristics of the application program, such as the language it is written in and the size of the largest I/O area it can use. The input deck can contain only one PSBGEN statement.

The END Statement

One END statement is placed at the end of each PSB generation input deck. The END statement specifies the end of input statements to the assembler.

Detailed instructions for coding PSB statements and examples of PSBs are contained in of *IMS Version 9: Utilities Reference: System*.

Building the Application Control Blocks (ACBGEN)

IMS builds the ACB with the ACBGEN utility by merging information from the PSB and DBD. For execution in a batch environment, IMS can build ACBs either dynamically (PARM=DLI), or it can prebuild them using the ACB Maintenance utility (PARM=DBB). ACBs *must* be prebuilt for use by online application programs. The ACB generation process is shown in Figure 183.

ACBs cannot be prebuilt for GSAM DBDs. However, ACBs can be prebuilt for PSBs that reference GSAM databases.

The ACB Maintenance utility, shown in Figure 183, gets the PSB and DBD information it needs from IMS.PSBLIB and IMS.DBDLIB.

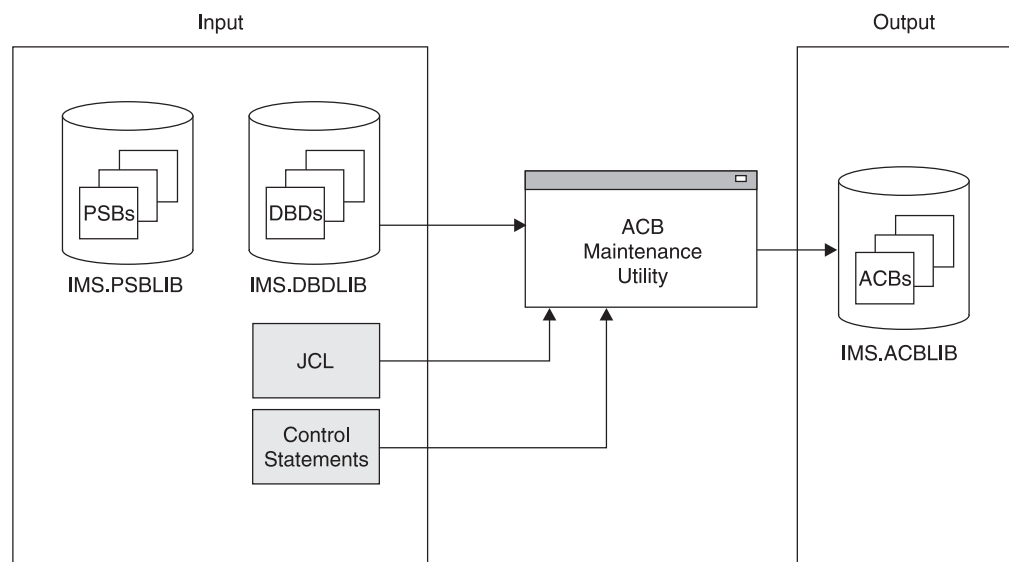


Figure 183. The ACB Generation Process

You can have the utility prebuild ACBs for all PSBs in IMS.PSBLIB, for a specific PSB, or for all PSBs that reference a particular DBD. Prebuilt ACBs are kept in the IMS.ACBLIB library. (IMS.ACBLIB is not used if ACBs are not prebuilt.) When ACBs are prebuilt and an application program is scheduled, the application program's ACB is read from IMS.ACBLIB directly into storage. This means that less time is required to schedule an application program. In addition, less storage is used if prebuilt ACBs are used. Another advantage of using the ACB Maintenance utility is the initial error checking it performs. It checks for errors in the names used in the PSB and the DBDs associated with the PSB and, if erroneous cross-references are found, prints appropriate error messages.

IMS.ACBLIB has to be used exclusively. Because of this, the ACB Maintenance utility can only be executed using an IMS.ACBLIB that is not currently allocated to an active IMS system. Also, because IMS.ACBLIB is modified, it cannot be used for any other purpose during execution of the ACB Maintenance utility.

You can change ACBs or add ACBs in an "inactive" copy of ACBLIB and then make the changed or new members available to an active IMS online system by using the

online change function. “Using the Online Change Function” in Chapter 16, “Modifying Databases,” on page 429 describes how you effectively change ACBLIB for an online system.

Detailed instructions for running the ACB Maintenance utility and examples of its use are contained in the *IMS Version 9: Utilities Reference: System*.

Defining Generated Program Specification Blocks for SQL Applications

Generated PSBs (GPSB) are a type of PSB that do not require a PSBGEN or ACBGEN. A GPSB contains an I/O PCB and a single modifiable alternate PCB. GPSBs are not defined through a PSBGEN. Instead, they are defined by the system definition process through the APPLCTN macro. The GPSB parameter indicates the use of a generated PSB and specifies the name to be associated with it. The LANG parameter specifies the language format of the GPSB. For more information on defining GPSBs refer to the APPLCTN macro topic of the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

The I/O PCB can be used by the application program to obtain input messages and send output to the inputting terminal. The alternate PCB can be used by the application program to send output to other terminals or programs.

Other than the I/O PCB, an application that makes only Structured Query Language (SQL) calls does not require any PCBs. It does, however, need to define the application program name and language type to IMS. A GPSB can be used for this purpose.

Chapter 12. Developing Test Databases

Before the application programs accessing your database are transferred to production status, they must be tested. To avoid damaging a production database, you need a test database.

IBM provides various programs that can help you develop your test database, including the DL/I Test Program (DFSDDLTO). For more information on the available IMS tools, go to www.ibm.com/ims and link to the IBM® DB2 and IMS Tools Web site.

Related Reading:

- For guidance information about application program testing, see *IMS Version 9: Application Programming: Design Guide*.
- For information about testing an online system, see *IMS Version 9: Administration Guide: System*.

In this Chapter:

“Test Requirements”

“Designing, Creating, and Loading a Test Database” on page 312

Test Requirements

Depending on your system configuration, user requirements, and the design characteristics of your database and data communication systems, test for the following:

- That DL/I call sequences execute and the results are correct.
 - This kind of test often requires only a few records, and you can use the DL/I Test Program, DFSDDLTO, to produce these records.
 - If this is part of a unit test, consider extracting records from your existing database. To extract the necessary records, you can use programs such as the IMS DataRefresher™.
- That calls execute through all possible application decision paths.
 - You might need to approximate your production database. To do this, you can use programs such as the IMS DataRefresher and other IMS tools.
- How performance compares with that of a model, for system test or regression tests, for example.
 - For this kind of test, you might need a copy of a subset of the production database. You can use IMS tools to help you.

To test for these capabilities, you need a test database that approximates, as closely as possible, the production database. To design such a test database, you should understand the requirements of the database, the sample data, and the application programs.

To protect your production databases, consider providing the test JCL procedures to those who test application programs. Providing the test JCL helps ensure that the correct libraries are used.

What Kind of Database?

Often, the test database can be a copy of a subset of the production database, or in some other way, a replica of it. If you have designed the production database, you should have firsthand knowledge of this requirement. Your DBDs for the production database can provide the details. If you have your production database defined in a data dictionary, that definition gives you much of the information you need. The topics in this chapter describe some aids available to help you design and generate your test database.

What Kind of Sample Data?

It is important for the sample data to approximate the real data, because you must test that the system processes data with the same characteristics, such as the range of field values. The kind of sample data needed depends on whether you are testing calls or program logic.

- To test calls, you need values in only those fields that are sequence fields or which are referenced in SSAs.
- To test program logic, you need data in all fields accessed in the program logic such as adds or compares.

Again, you might use a copy of a subset of the real database. However, first determine which fields contain sensitive data and therefore must use fictitious data in the test database.

What Kind of Application Program?

In order to design a test database that effectively tests the operational application programs being developed, you should know certain things about those programs. Much of the information you need is in the application program design documentation, the descriptors such as the PSBs, your project test plan, and in the Data Dictionary.

Designing, Creating, and Loading a Test Database

You can develop a test database just as you would develop a production database. With that approach, you perform the tasks described throughout the other chapters of this manual, keeping in mind the special requirements for test databases. If your installation has testing standards and procedures, you should follow them in developing a test database.

Using Testing Standards

Testing standards and procedures help you avoid the same kinds of problems for test database development as your IMS development standards do for production databases. Some of the subjects that might be included in your test system standards and that affect test database design are:

- Objectives of your test system
 - What you test for and at what development stages do you test for it
 - The kinds of testing—offline, online, integrated DB/DC or isolated
- Description of the test organization and definition of responsibilities of each group
- Relationship of test and production modes of operation
- How your test system development process deals with:
 - DB/TM structures
 - Development tools

- DB/TM features
- Backup and recovery

Related Reading: The IMS test system is discussed in *IMS Version 9: Administration Guide: System*.

Using IBM Programs to Develop a Test Database

If you use the same development aids to develop the test database that you use to develop your production databases, you will benefit from using familiar tools. Also, you will avoid problems caused by differences between test and production databases.

Using the Data Extraction, Processing, and Restructuring System

You can use this system (Program Number: 5796-PLH) to access a wide variety of data and restructure it into a test database. By means of control statements, you define the source and target files and specify the structure of the target files.

The data restructuring phase of the system allows you to:

- Combine components of different files to form new files
- Restructure a file to form different files
- Rearrange data within a file
- Alter values according to your needs
- Form hierarchies
- Decrease or increase the number of levels in a hierarchy

Details about using this system are in *Data Extraction, Processing, and Restructuring System, Program Description/Operations Manual*.

Using IBM VisualAge Generator

IBM VisualAge® Generator is the direct descendant of IBM's Cross System Product mainframe application development tool. VisualAge Generator provides a rapid application development environment for building and deploying multi-tier and text applications and applications for e-business. VisualAge Generator allows developers to implement end-to-end Java systems for e-business and enables OO and Java developers with little mainframe experience to implement systems on CICS and IMS transactional platforms that access DL/I or VSAM data.

VisualAge Generator systems can reuse existing legacy application programs, and VisualAge Generator application programs can access VSAM file systems and IMS and DL/I databases. This allows easy integration of new applications programs into existing IT infrastructures.

Related Reading: For information on how to use VisualAge Generator, see the *VisualAge Generator: User's Guide*.

VisualAge Generator supports both IMS and CICS.

IMS Application Development Facility II

The IMS Application Development Facility II (IMSADF II) is designed to increase the productivity of programmers developing applications for IMS systems. IMSADF II provides ISPF-based dialogs to work with IMS and DB2 to reduce the time and effort that is required to develop, maintain, and expand IMS database and data communications applications.

Designing, Creating, and Loading a Test Database

If your installation uses IMSADF II to develop application programs, you can use it to create a simple test database. The interactive nature of the IMSADF II enables you to dynamically add segments to a database. By means of SEGM and FIELD statements, you can define a test database and update it as needed.

Related Reading: For information on how to use the IMS Application Development Facility II, see the *IMS Application Development Facility (ADF) II: User's Guide*.

File Manager for z/OS for IMS Data

The IMS component of File Manager (FM/IMS) is an ISPF application with which you can manipulate data stored in IMS databases.

FM/IMS provides you with a number of flexible ways to connect to your IMS databases. For example, with BMP mode you can connect to an online multi-user database and manipulate the data. In DLI mode, you can work with data offline as a single user or you can share the data with others.

In addition, FM/IMS provides two functions that you can use in batch jobs. FM/IMS Edit Batch (IEB) runs a REXX procedure that can insert, update, retrieve, delete or print segments and create views. FM/IMS Batch Print (IPR) can print the entire database in one of several available display formats, or a selected subset of the database, based on a view.

Related Reading: For information on how to use the File Manager for z/OS for IMS Data, see the *File Manager for z/OS User's Guide and Reference for IMS Data*.

Using the DL/I Test Program, DFSDDLTO

If you need a test database with relatively few database records, for example, you can use DFSDDLTO to test DL/I call sequences. If you have no machine-readable database to begin with, you can define a PCB, then use DFSDDLTO to insert segments. This step eliminates the need for a load program to generate your test database.

Related Reading: Information about this test program is in "Testing an Application Program," in *IMS Version 9: Application Programming: Design Guide*.

The DL/I Test Program cannot be used by CICS, but can be used for stand-alone batch programs. If used for stand-alone batch programs, it is useful to interpret the database performance as it might be implemented for online or shared database programs.

Chapter 13. Loading Databases

After you implement your database design, you are ready to write and load your database. However, before writing a load program, you must estimate the minimum size of the database and allocate data sets.

This chapter examines the following areas of loading a database:

- “Estimating the Minimum Size of the Database”
- “Allocating Data Sets” on page 322
- “Writing a Load Program” on page 325
- “Loading Fast Path Databases” on page 337

Estimating the Minimum Size of the Database

When you estimate the size of your database, you estimate how much space you need to initially load your data. Unless you do not plan to insert segments into your database after it is loaded, allocate more space for your database than you actually estimate for the initial load.

This topic contains the step-by-step procedure for estimating minimum database space. To estimate the minimum size needed for your database, you must already have made certain design decisions about the physical implementation of your database. Because these decisions are different for each DL/I access method, they are discussed under the appropriate access method in step 3 of the procedure.

If you plan to reorganize your HALDBs online, consider the extra space reorganization requires. Although online reorganization does not need any additional space when you first load a HALDB, the process does require additional space at the time of reorganization. For more information on HALDB online reorganization, see “HALDB Online Reorganization” on page 370.

Step 1. Calculate the Size of an Average Database Record

First, determine the size, then the average number of occurrences of each segment type in a database record. By multiplying these two numbers together, you get the size of an average database record.

Determining Segment Size

Segment size here is physical segment size, and it includes both the prefix and data portion of the segment. You define the size of the data portion. It can include unused space for future use. The size of the data portion of the segment is the number you specified in the BYTES= operand in the SEGM statement in the DBD.

The prefix portion of the segment depends on the segment type and on the options you are using. Table 23 on page 316 helps you determine, by segment type, the size of the prefix. Using the chart, add up the number of bytes required for necessary prefix information and for extra fields and pointers generated in the prefix for the options you have chosen. Segments can have more than one 4-byte pointer in their prefix. You need to factor all extra pointers of this type into your calculations.

Related Reading: For rules on using mixed pointers, see “Mixing Pointers” on page 93.

Estimating the Minimum Size of the Database

Table 23. Required Fields and Pointers in a Segment's Prefix

| Type of Segment | Fields and Pointers Used in the Segment's Prefix | Size of the Field or Pointer (in Bytes) |
|---------------------------------------|---|--|
| All types | Segment code (not present in a SHSAM, SHISAM, GSAM, or secondary index pointer segment) | 1 |
| | Delete byte (not present in a SHSAM, SHISAM, or GSAM segment) | 1 |
| HDAM, PHDAM, HIDAM, and PHIDAM | PCF pointer | 4 |
| | PCL pointer | 4 |
| | PP pointer | 4 |
| | PTF pointer | 4 |
| | PTB pointer | 4 |
| HDAM and HIDAM only | HF pointer | 4 |
| | HB pointer | 4 |
| DEDB | PCF pointer | 4 |
| | PCL pointer | 4 |
| | Subset pointer | 4 |
| Logical parent (for HDAM and HIDAM) | LCF pointer | 4 |
| | LCL pointer | 4 |
| | Logical child counter | 4 |
| Logical parent (for PHDAM and PHIDAM) | Logical child counter (only present for unidirectional logical parents) | 4 |
| Logical child | LTF pointer | 4 |
| | LTB pointer | 4 |
| | LP pointer | 4 |
| Logical child (PHDAM and PHIDAM) | EPS | 28 |
| Secondary index | Symbolic or direct-address pointer to the target segment | 4 |
| PSINDEX | EPS plus the target segment root key | 28 + length of the target-segment root key |
| All segments in PHDAM and PHIDAM | ILK | 8 |

Determining Segment Frequency

After you have determined the total size of a segment type, you need to determine segment frequency. Segment frequency is the average number of occurrences of a particular segment type in the database record. To determine segment frequency, first determine the average number of times a segment occurs under its immediate physical parent.

For example, in the database record in Figure 184 on page 317, the ITEMS segment occurs an average of 10 times for each DEPOSITS segment. The DEPOSITS segment occurs an average of four times for each CUSTOMER root segment. The frequency of a root segment is always one.

Estimating the Minimum Size of the Database

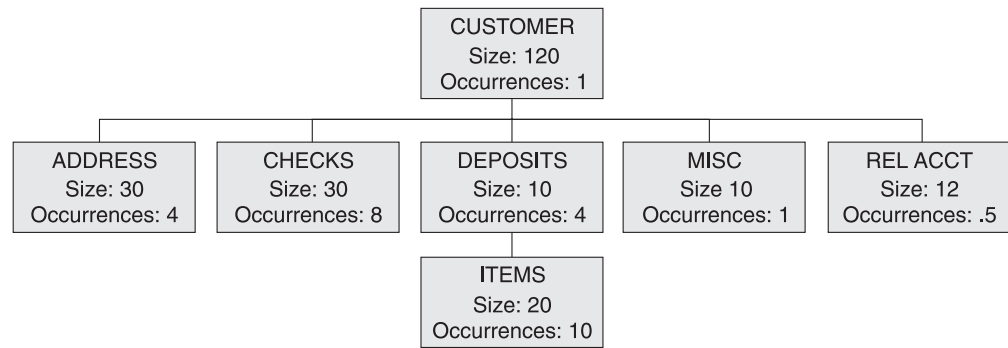


Figure 184. Segment Sizes and Average Segment Occurrences

To determine the average number of occurrences of a particular segment type in the database record, multiply together the segment frequencies of each segment in the path from the given segment back up to the root. For the ITEMS segment type, the path includes the ITEMS segment and the DEPOSITS segment. The segment frequency of ITEMS is 10, and the segment frequency of DEPOSITS is 4. Therefore, the average number of occurrences of the ITEMS segment in the database record is 40 (10 x 4). Another way of expressing this idea is that each customer has an average of 4 DEPOSITS, and each DEPOSIT has an average of 10 ITEMS. Therefore, for each customer, an average of 40 (10 x 4) ITEMS exist in the database record.

Determining Average Database Record Size

Now that you have determined segment size and segment frequency, you can determine the average size of a database record. To determine average database record size for a HISAM database, multiply segment size and segment frequency together for each segment type in the database record, then add the results. For example, for the database record in Figure 184, the average database record size is calculated as shown in Table 24.

Table 24. Calculating the Average Database Record Size

| Segment Type | Segment Size | Average Occurrences | Total Size |
|---------------------|--------------|---------------------|-------------|
| CUSTOMER | 120 | 1 | 120 |
| ADDRESS | 30 | 4 | 120 |
| CHECKS | 30 | 8 | 240 |
| DEPOSITS | 10 | 4 | 40 |
| ITEMS | 20 | 40 (10x4) | 800 |
| MISC | 10 | 1 | 10 |
| REL ACCT | 12 | .5 | 6 |
| Record Total | | | 1336 |

Step 2. Determine Overhead Needed for CI Resources

If you are not using VSAM, you can skip this step. If you are using VSAM, you need to determine how much overhead is needed for a CI before you can do the remaining space calculations.

Overhead is space used in a CI for two control fields. VSAM uses the control fields to manage space in the CI. The control fields and their sizes are shown in Table 25 on page 318

Estimating the Minimum Size of the Database

on page 318.

Table 25. VSAM Control Fields

| Field | Size in Bytes |
|--|---------------|
| CIDF (Control interval definition field) | 4 |
| RDF (Record definition field) | 3 |

If one logical record exists for each CI, CI overhead consists of one CIDF and one RDF (for a total of 7 bytes). HDAM and HIDAM databases and PHDAM and PHIDAM partitions use one logical record for each CI.

If more than one logical record exists for each CI, CI overhead consists of one CIDF and two RDFs (for a total of 10 bytes). HISAM (KSDS and ESDS), HIDAM and PHIDAM index, and secondary index databases can all use more than one logical record for each CI.

Step 3 tells you when to factor CI overhead into your space calculations.

Step 3. Determine the Number of CIs or Blocks Needed

The calculations in this step are done by database type. To determine how many CIs or blocks are needed to hold your database records, go to the topic in this step that applies to the database type you are using. If you are using VSAM, the first CI in the database is reserved for VSAM.

HISAM: Determining the Number of CIs or Blocks Needed

A CI in HISAM can contain one or more logical records. In the primary data set a logical record can only contain one database record (or part of one database record). In the overflow data set a logical record can only contain segments of the same database record, but more than one logical record can be used for the overflow segments of a single database record.

In HISAM, you should remember how logical records work, because you need to factor logical record overhead into your calculations before you can determine how many CIs (control intervals) are needed to hold your database records. Logical record overhead is a combination of the overhead that is always required for a logical record and the overhead that exists because of the way in which database records are stored in logical records (that is, storage of segments almost always results in residual or unused space).

Because some overhead is associated with each logical record, you need to calculate the amount of space that is available after factoring in logical record overhead. Once you know the amount of space in a logical record available for data, you can determine how many logical records are needed to hold your database records. If you know how many logical records are required, you can determine how many CIs or blocks are needed.

For example, assume you need to load 500 database records using VSAM, and to use a CI size of 2048 bytes for both the KSDS and ESDS. Also, assume you need to store four logical records in each KSDS CI and two logical records in each ESDS CI.

1. First factor in CI overhead by subtracting the overhead from the CI size: $2048 - 10 = 2038$ bytes for both the KSDS and the ESDS. The 10 bytes of overhead consists of a 4-byte CIDF and two 3-byte RDFs.

2. Then, calculate logical record size by dividing the available CI space by the number of logical records per CI: $2038/4 = 509$ bytes for the KSDS and $2038/2 = 1019$ bytes for the ESDS. Because logical record size in HISAM must be an even value, use 508 bytes for the KSDS and 1018 bytes for the ESDS.
3. Finally, factor in logical record overhead by subtracting the overhead from logical record size: $508 - 5 = 503$ bytes for the KSDS and $1018 - 5$ bytes for the ESDS. HISAM logical record overhead consists of 5 bytes for VSAM (a 4-byte RBA pointer for chaining logical records and a 1-byte end-of-data indicator).
This means if you specify a logical record size of 508 bytes for the KSDS, you have 503 bytes available in it for storing data. If you specify a logical record size of 1018 bytes for the ESDS, you have 1013 bytes available in it for storing data.

Refer to the previous example. Because the average size of a database record is 1336 bytes, the space available for data in the KSDS is not large enough to contain it. It takes the available space in one KSDS logical record plus one ESDS logical record to hold the average database record ($503 + 1013 = 1516$ bytes of available space). This record size is greater than an average database record of 1336 bytes. Because you need to load 500 database records, you need 500 logical records in both the KSDS and ESDS.

- To store four logical records per CI in the KSDS, you need a minimum of $500/4 = 125$ CIs of 2048 bytes each for the KSDS.
- To store two logical records per CI in the ESDS, you need a minimum of $500/2 = 250$ CIs of 2048 bytes each for the ESDS.

HIDAM or PHIDAM: Determining the Number of CIs or Blocks Needed

With HIDAM or PHIDAM, one VSAM logical record exists per CI or block. In this context, logical record is the unit of transfer when invoking an access method (such as VSAM), to get or put records. Logical record overhead consists of an FSEAP (4 bytes). If you are using RAPs (HIDAM only), the logical record overhead consists of one RAP (4 bytes). For example, assume you need to load 500 database records using VSAM and to use a CI size of 2048 bytes and no RAP (specify PTR=TB on the root to suppress the RAP for HIDAM).

1. First, determine the size of a logical record by subtracting CI overhead from CI size: $2048 - 7 = 2041$ bytes for the ESDS logical record size. The 7 bytes of overhead consists of a 4-byte CIDE and a 3-byte RDF.
2. Then, determine the amount of logical record space available for data by factoring in logical record overhead. In this example, logical record overhead consists of an FSEAP: $2041 - 4 = 2037$ bytes. This means you have 2037 bytes available to store data in each logical record.

HIDAM or PHIDAM Index: Calculating the Space Needed

Calculating space for a HIDAM or PHIDAM index is similar to calculating space for a HISAM KSDS. The difference is that no logical record overhead exists. One index record is stored in one logical record, and multiple logical records can be stored in one CI or block.

HDAM or PHDAM: Determining the Amount of Space Needed

Because of the many variables in HDAM or PHDAM, no exact formula exists for estimating database space requirements. Therefore, you should use a space calculation aid to help determine the amount of space needed for HDAM or PHDAM databases.

Estimating the Minimum Size of the Database

If you are using VSAM, and you decide to estimate, without use of an aid, the amount of space to allocate for the database, the first CI in the database is reserved for VSAM. Because of this, the bit map is in the second CI.

With HDAM or PHDAM, logical record overhead depends on the database design options you have selected. You must choose the number of CIs or blocks in the root addressable area and the number of RAPS for each CI or block. These choices are based on your knowledge of the database.

A perfect randomizer requires as many RAPS as there are database records. Because a perfect randomizer does not exist, plan for approximately 20% more RAPS than you have database records. The extra RAPS reduces the likelihood of synonym chains. For example, assume you need to store 500 database records. Then, for the root addressable area, if you use:

- One RAP per CI or block, you need 600 CIs or blocks
- Two RAPs per CI or block, you need 300 CIs or blocks
- Three RAPs per CI or block, you need 200 CIs or blocks

Because of the way your randomizer works, you decide 300 CIs or blocks with two RAPs each works best. Assume you need to store 500 database records using VSAM, and you have chosen to use 300 CIs in the root addressable area and two RAPs for each CI. This decision influences your choice of CI size. Because you are using two RAPs per CI, you expect two database records to be stored in each CI. You know that a 2048-byte CI is not large enough to hold two database records ($2 \times 1336 = 2672$ bytes). And you know that a 3072-byte CI is too large for two database records of average size. Therefore, you would probably use 2048-byte CIs and the byte limit count to ensure that on average you would store two database records in the CI.

To determine the byte limit count:

1. First, determine the size of a logical record by subtracting CI overhead from CI size: $2048 - 7 = 2041$ bytes for the ESDS logical record size.
2. Then, determine the amount of logical record space available for data by factoring in logical record overhead. (Remember only one logical record exists per CI in HDAM or PHDAM.) In this example, logical record overhead consists of a 4-byte FSEAP and two 4-byte RAPs: $2041 - 4 - (2 \times 4) = 2029$ bytes. This means you have 2029 bytes available for storing data in each logical record in the root addressable area.
3. Finally, determine the available space per RAP by dividing the available logical record space by the number of RAPs per CI: $2029/2 = 1014$ bytes. Therefore, you must use a byte limit count of about 1000 bytes.

Continuing our example, you know you need 300 CIs of 2048 bytes each in the root addressable area. Now you need to calculate how many CIs you need in the overflow area. To do this:

- Determine the average number of bytes that will not fit in the root addressable area. Assume a byte limit count of 1000 bytes. Subtract the byte limit count from the average database record size: $1336 - 1000 = 336$ bytes. Multiply the average number of overflow bytes by the number of database records: $500 \times 336 = 168000$ bytes needed in the non-root addressable area.
- Determine the number of CIs needed in the non-root addressable area by dividing the number of overflow bytes by the bytes in a CI available for data. Determine the bytes in a CI available for data by subtracting CI and logical

record overhead from CI size: $2048 - 7 - 4 = 2037$ (7 bytes of CI overhead and 4 bytes for the FSEAP). Overflow bytes divided by CI data bytes is $168000/2037 = 83$ CIs for the overflow area.

You have estimated you need a minimum of 300 CIs in the root addressable area and a minimum of 83 CIs in the non-root addressable area.

Secondary Index: Determining the Amount of Space Needed

Calculating space for a secondary index is similar to calculating space for a HISAM KSDS. The difference is that no logical record overhead exists in which factor. One index record is stored in one logical record, and multiple logical records can be stored in one CI or block.

Step 4. Determine the Number of Blocks or CIs Needed for Free Space

In HDAM, HIDAM, PHDAM, and PHIDAM databases, you can allocate free space when your database is initially loaded. Free space is explained in Chapter 6, “Choosing Full-Function Database Types,” on page 59, “Specifying Free Space”. Free space can only be allocated for an HD VSAM ESDS or OSAM data set. Do not confuse the free space discussed here with the free space you can allocate for a VSAM KSDS using the DEFINE CLUSTER command.

To calculate the total number of CIs or blocks you need to allocate in the database, you can use the following formula:

$$A = B \times (\text{fbff} / (\text{fbff} - 1)) \times (100 / (100 - \text{fspf}))$$

Where the values are:

- A** The total number of CIs or blocks needed including free space.
- B** The number of blocks or CIs in your database.
- fbff** How often you are leaving a block or CI in the database empty for free space (what you specified in fbff operand in the DBD).
- fspf** The minimum percentage of each block or CI you are leaving as free space (what you specified in the fspf operand in the DBD).

Step 5. Determine the Amount of Space Needed for Bit Maps

In HDAM, HIDAM, PHDAM, and PHIDAM databases, you need to add the amount of space required for bit maps to your calculations. Bit maps are explained in “General Format of HD Databases and Use of Special Fields” on page 95. To calculate the number of bytes needed for bit maps in your database, you can use the following formula:

$$A = D / ((B - C) \times 8)$$

Where the values are:

- A** The number of bit map blocks or CIs you need for the database.
- B** The CI or block size you have specified, in bytes, minus 4.
Four is subtracted from the CI or block size because each CI or block has a 4-byte FSEAP.
- C** The number of RAPs you specified for a CI or block, times 4.
The number of RAPs is multiplied by 4 because each RAP is four bytes long. $(B - C)$ is multiplied by 8 in the formula to arrive at the total number of bits that will be available in the CI or block for the bit map.

Estimating the Minimum Size of the Database

D The number of CIs or blocks in your database.

You need to add the number of CIs or blocks needed for bit maps to your space calculations.

Allocating Data Sets

Once you have determined how much space you will need for your database, you can allocate data sets and then load your database.

You can allocate VSAM data sets by using the `DEFINE CLUSTER` command. You must specify the `REUSE` parameter when allocating HALDB data sets. Use of this command is described in *DFSMS Access Method Services for Catalogs*.

Attention: If you plan to use the Database Image Copy 2 utility to take image copies of your database, the data sets must be allocated on hardware that supports the DFSMS concurrent copy function.

When loading databases (excluding HALDB databases) that contain logical relationships or secondary indexes, DL/I writes a control record to a work file (DFSURWF1). This work file must also be allocated and in the JCL.

All other data sets are allocated using normal z/OS JCL. You can use the z/OS program IEFBR14 to preallocate data sets, except when the database is an MSDB. For MSDBs, you should use the z/OS program IEHPROGM.

Allocating OSAM Data Sets

For databases other than HALDBs, at the time the data set is loaded, you should use JCL to allocate OSAM data sets. For HALDB OSAM data sets, the allocation must be done before the load. This mode of allocation can be for single or multiple volumes, using the `SPACE` parameter.

If the installation control of direct-access storage space and volumes require that the OSAM data sets be pre-allocated, or if a message queue data set requires more than one volume, the OSAM data sets might be pre-allocated.

Observe the following restrictions when you preallocate with any of the accepted methods:

- DCB parameters should not be specified.
- Secondary allocation must be specified for all volumes if the data set will be extended beyond the primary allocation.
- Secondary allocation must be specified for all volumes in order to write to volumes pre-allocated but not written to by initial load or reload processing.
- Secondary allocation is not allowed for queue data sets because queue data sets are not extended beyond their initial or pre-allocated space quantity. However, queue data sets can have multivolume allocation.
- The secondary allocation size defined on the first volume will be used for all secondary allocations on all volumes regardless of the secondary allocation size specified on the other volumes. All volumes should be defined with the same secondary allocation size to avoid confusion.
- If the OSAM data set will be cataloged, use IEHPROGM or Access Method Services to ensure that all volumes are included in the catalog entry.

When a multiple-volume data set is pre-allocated, you should allocate extents on all the volumes to be used. The suggested method of allocation is to have one IEFBR14 utility step for each volume on which space is desired.

Restrictions:

- **Do not** use IEFBR14 and specify a DD card with a multivolume data set, because this allocates an extent on only the first volume.
- **Do not** use a separate IEFBR14 utility step for each volume if you plan to use the Database Image Copy 2 utility (DFSUDMT0). Using a separate IEFBR14 step for each volume creates extents that have the same identifier (volume 1). The Database Image Copy 2 utility requires that you allocate your multi-volume data sets using the standard DFSMS methods.

For more information about the standard DFSMS methods for allocating data sets and about data sets in general, see:

- *z/OS: DFSMS: Using Data Sets*
- *z/OS: DFSMS Access Method Services for Catalogs*

Example of Allocating Multi-Volume OSAM Data Sets on Non-SMS Managed DASD

The JCL in Figure 185 is an example of allocating a multi-volume OSAM data set on a non-SMS managed direct-access storage device (DASD). The data sets allocated with this JCL are not compatible with the Database Image Copy 2 utility.

```
//OSAMALLO JOB A,OSAMEXAMPLE
//S1 EXEC PGM=IEFBR14
//SYSPRINT DD SYSOUT=A
//EXTENT1 DD VOL=SER=AAAAAA,SPACE=(CYL,(20,5)),UNIT=3390,
// DSN=OSAM.SPACE,DISP=(,KEEP)
//S2 EXEC PGM=IEFBR14
//SYSPRINT DD SYSOUT=A
//EXTENT2 DD VOL=SER=BBBBBB,SPACE=(CYL,(30,5)),UNIT=3390,
// DSN=OSAM.SPACE,DISP=(,KEEP)
.
.
.
//LAST EXEC PGM=IEFBR14
//SYSPRINT DD SYSOUT=A
//EXTENTL DD VOL=SER=LLLLLL,SPACE=(CYL,(30,5)),UNIT=3390,
// DSN=OSAM.SPACE,DISP=(,KEEP)
// EXEC PGM=IEHPROGM
//SYSPRINT DD SYSOUT=*
//DD1 DD UNIT=3390,VOL=SER=AAAAAA,DISP=SHR
//DD2 DD UNIT=3390,VOL=SER=BBBBBB,DISP=SHR
.
.
.
//DDL DD UNIT=3390,VOL=SER=LLLLLL,DISP=SHR
//SYSIN DD *
CATLG DSNAME=OSAM.SPACE,VOL=3390=(AAAAAA,BBBBBB,LLLLLL)
/*
```

Figure 185. JCL allocating a multi-volume OSAM data set on non-SMS managed DASD

Example of Allocating Multi-Volume OSAM Data Sets on SMS-Managed DASD

Figure 186 is an example of the JCL you can use to allocate multi-volume OSAM data sets that are compatible with the Database Image Copy 2 utility.

When allocating data sets using the method in Figure 186, you must specify an SMS storage class with the guaranteed space attribute, GTDSPACE, and specify the volume serial numbers. Without the guaranteed space attribute and the volume serial numbers, only the first volume will get a primary extent and the other volumes will get secondary extents.

```
//          EXEC   PGM=IEFBR14
//DD123      DD   DSN=HIGHNODE.MIDNODE.ENDNODE,
//           DISP=(NEW,CATLG),
//           SPACE=(CYL,(200,100)),
//           UNIT=3390
//           VOL=SER=(VOL111,VOL222,VOL333),
//           STORCLAS=GTDSPACE
/*
```

Figure 186. Example JCL for Allocating a Multi-Volume OSAM Data Set on SMS-Managed DASD

Figure 187 is an example of using DFSMS Access Method Services ALLOCATE command to allocate multi-volume OSAM data sets that are compatible with the Database Image Copy 2 utility.

```
//          EXEC   PGM=IDCAMS,REGION=4096K
//SYSPRINT   DD   SYSOUT=A
//SYSIN      DD   *
              DELETE 'HIGHNODE.MIDNODE.ENDNODE'

              IF LASTCC<=8 THEN SET MAXCC=0

              ALLOC DSN('HIGHNODE.MIDNODE.ENDNODE')-
              NEW CATALOG-
              SPACE(200,100)CYLINDERS-
              UNIT=3390-
              VOL(VOL111,VOL222,VOL333)-
              STORCLAS(GTDSPACE)
/*
```

Figure 187. Example of ALLOCATE Command for Allocating a Multi-Volume OSAM Data Set on SMS-Managed DASD

Cautions When Allocating OSAM Data Sets

1. Pre-allocating more volumes for OSAM data set extents than are used during initial load or reload processing might cause an abend if you attempt to extend the data set beyond the last volume written to at initial load or reload time under the following circumstances: the initial load or reload step did not result in the data being written to the last volume of the pre-allocated data set, and secondary allocation was not specified during data set pre-allocation.
2. It is recommended that you not reuse multivolume OSAM data sets without first scratching the data set and then reallocating the space. Failure to do this might cause an invalid EOF mark to be left in the DSCB of the last volume of the data set when the data set is:
 - a. First reused by an IMS utility (such as the Unload/Reload utility used in database reorganization).

- b. Then opened by OSAM for normal processing.

For example, a data set might initially be allocated on three volumes, with the EOF mark on the third volume. However, after the reorganization utility is run, the data set might need only the first two volumes. Therefore, the new EOF mark is placed on the second volume. After reorganization, when the data set is opened by OSAM for normal processing, OSAM checks the last volume's DSCB for an EOF mark. When OSAM finds the EOF in the third volume, it inserts new data after the old EOF mark in the third volume instead of inserting data after the EOF mark created by the reorganization utility in the second volume.

Subsequent processing by another utility such as the Image Copy utility uses the EOF mark set by the reorganization utility on the second volume and ignores new data inserted by OSAM on volume three.

3. When loading this database, the order of the DD cards determines the order in which the data is loaded.
4. If you intend to use the Image Copy 2 utility (DFSUDMT0) to back up and restore multi-volume databases, they **MUST** be allocated using the standard DFSMS techniques.

For more information about the DFSMS methods of data set allocation, see *thez/OS V1R4: DFSMS: Using Data Sets*.

Writing a Load Program

After you have determined how much space your database requires and allocated data sets for it, you can load the database.

The Load Process

Loading the database is done using an initial load program. Initial load programs must be batch programs, since you cannot load a database with an online application program. It is your responsibility to write this program.

Basically, an initial load program reads an existing file containing your database records. Using the DBD, which defines the physical characteristics of the database, and the load PSBs (see Figure 188 on page 328), the load program builds segments for a database record and inserts them into the database in hierarchic order. If the data to be loaded into the database already exists in one or more files (see Figure 189 on page 329), merge and sort the data, if necessary, so that it is presented to the load program in correct sequence. Also, if you plan to merge existing files containing redundant data into one database, delete the redundant data, if necessary, and correct any data that is wrong.

After you have defined the database, you load it by writing an application program that uses the ISRT call. An initial load program builds each segment in the program's I/O area, then loads it into the database by issuing an ISRT call for it. ISRT calls are the only DL/I requests allowed when you specify PROCOPT=L in the PCB. The only time you use the "L" option is when you initially load a database. This option is valid only for batch programs.

Recommendation: If a user load program using PROCOPT=LILS is running in a DLI or DBB region, DBRC authorization is required for all databases logically related to the one being loaded. If DBRC is active when the database is loaded, DBRC sets the image copy status for this database to IC NEEDED in the DBDS record in the RECON data set.

Writing a Load Program

The FIRST, LAST, and HERE insert rules do not apply when you are loading a database, unless you are loading an HDAM database. When you are loading a HDAM database, the rules determine how root segments with non-unique sequence fields are ordered. If you are loading a database using HSAM, the same rules apply.

Recommendation: Load programs do not need to issue checkpoints.

Most comprehensive databases are loaded in stages by segment type or by groups of segment types. Because there are usually too many segments to load using only one application program, you need several programs to do the loading. Each load program after the first load program is technically an “add” program, not a load program. Do not specify “L” as the processing option in the PCB for add programs. You should review any add type of load program written to load a database to ensure that the program’s performance will be acceptable; it usually takes longer to add a group of segments than to load them.

For HSAM, HISAM, HIDAM, and PHIDAM, the root segments that the application program inserts must be pre-sorted by the key fields of the root segments. The dependents of each root segment must follow the root segment in hierarchic sequence, and must follow key values within segment types. In other words, you insert the segments in the same sequence in which your program would retrieve them if it retrieved in hierarchic sequence (children after their parents, database records in order of their key fields).

If you are loading an HDAM or PHDAM database, you do not need to pre-sort root segments by their key fields.

When you load a database:

- If a loaded segment has a key, the key value must be in the correct location in the I/O area.
- When you load a logical child segment, the I/O area must contain the logical parent’s concatenated key, followed by the logical child segment to be inserted.
- After issuing an ISRT call, the current position is just before the next available space following the last segment successfully loaded. The next segment you load will be placed in that space.

Recommendation: You should always create an image copy immediately after you load, reload, or reorganize the database.

Status Codes for Load Programs

If the ISRT call is successful, DL/I returns a blank status code for the program. If not, DL/I returns one of these status codes:

LB The segment you are trying to load already exists in the database. DL/I only returns this status code for segments with key fields.

In a call-level program, you should transfer control to an error routine.

LC The segment you are trying to load is out of key sequence.

LD No parent exists for this segment. This status code usually means that the segment types you are loading are out of sequence.

LE In an ISRT call with multiple SSAs, the segments named in the SSAs are not in their correct hierarchic sequence.

- LF** Initial load of PHDAM or PHIDAM attempted ISRT of a logical child segment.
- V1** You have supplied a variable-length segment whose length is invalid.

Using SSAs in a Load Program

When you are loading segments into the database, you do not need to worry about position, because DL/I inserts one segment after another. The most important part of loading a database is the order in which you build and insert the segments.

The only SSA you must supply is the unqualified SSA giving the name of the segment type you are inserting.

Because you do not need to worry about position, you need not use SSAs for the parents of the segment you are inserting. If you do use them, be sure they contain only the equal (EQ, =b, or b=) relational operator. You must also use the key field of the segment as the comparative value.

For HISAM, HIDAM, and PHIDAM, the key X'FFFF' is reserved for IMS. IMS returns a status code of LB if you try to insert a segment with this key.

Loading a Sequence of Segments with the D Command Code

You can load a sequence of segments in one call by concatenating the segments in the I/O area and supplying DL/I with a list of unqualified SSAs. You must include the D command code with the first SSA. The sequence that the SSAs define must lead down the hierarchy, with each segment in the I/O area being the child of the previous segment.

Two Types of Initial Load Program

Two types of initial load programs exist: *basic* and *restartable*. The basic program must be restarted from the beginning if problems occur during execution. The restartable program can be restarted at the last checkpoint taken before problems occurred. Restartable load programs must be run under control of the Utility Control Facility (UCF) and require VSAM as the access method. The following topics describe both types of load programs:

- “Basic Initial Load Program” on page 329
- “Restartable Initial Load Program” on page 332

Figure 188 on page 328 shows the load process.

Figure 189 on page 329 illustrates loading a database using existing files.

Writing a Load Program

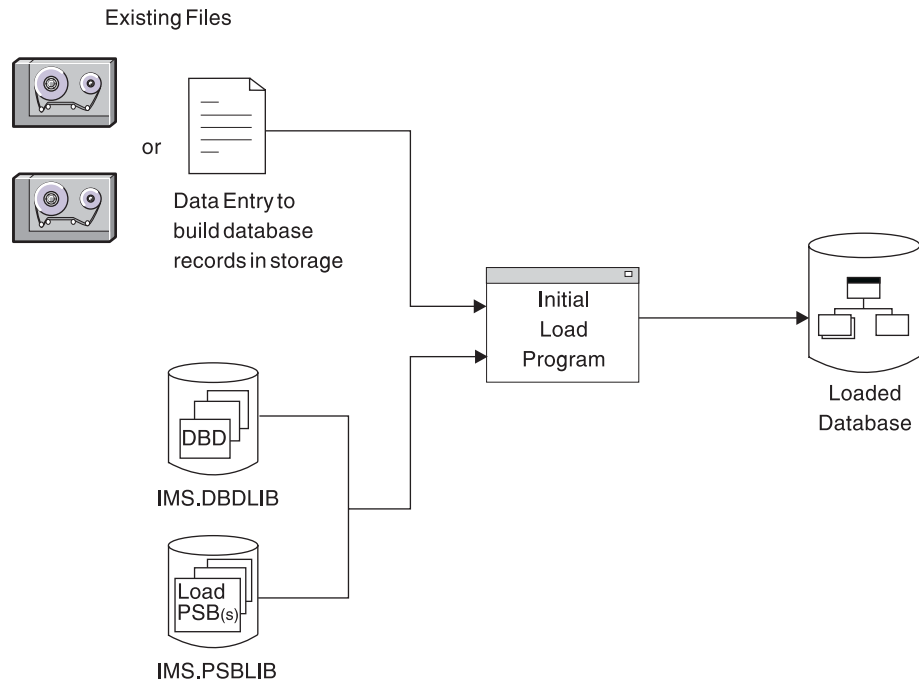


Figure 188. The Load Process

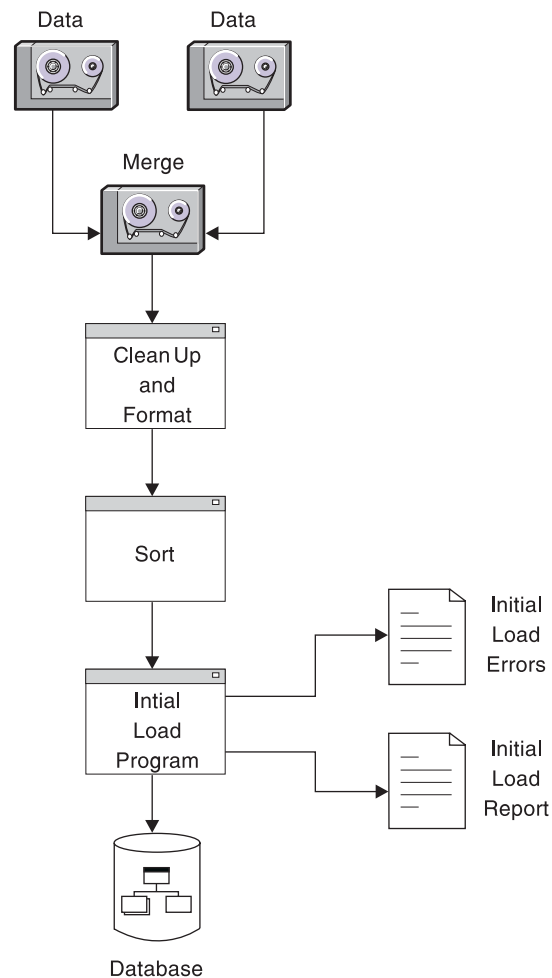


Figure 189. Loading a Database Using Existing Files

Basic Initial Load Program

You should write a basic initial load program (one that is not restartable) when the volume of data you need to load is not so great that you would be seriously set back if problems occurred during program execution. If problems do occur, the basic initial load program must be rerun from the beginning.

Figure 190 on page 331 shows the logic for developing a basic initial load program. Following Figure 190 is a sample load program (Figure 191) that satisfies the basic IMS database loading requirements. A sample program showing how this can be done with the Utility Control Facility is also provided.

Fast Path Data Entry Databases (DEDBs) cannot be loaded in a batch job as can DL/I databases. DEDBs are first initialized by the DEDB Initialization Utility and then loaded by a user-written Fast Path application program that executes typically in a BMP region.

Related Reading: See *IMS Version 9: Utilities Reference: Database and Transaction Manager* for a description of how DEDBs are loaded.

Writing a Load Program

Fast Path Main Storage Databases (MSDBs) are not loaded until the IMS control region is initialized. These databases are then loaded by the IMS start-up procedure when the following requirements are met:

- The MSDB= parameter on the EXEC Statement of Member Name IMS specifies a one-character suffix to DBFMSDB in IMS.PROCLIB.
- The member contains a record for each MSDB to be loaded.

The record contains a record for each MSDB, the number of segments to be loaded, and an optional “F” which indicates that the MSDB is to be fixed in storage.

Related Reading: For a description of the record format and the DBD keyword parameters, see the topics about member name IMS in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

- A sequential data set, part of a generation data group (GDG) with dsname IMS.MSDBINIT(0), is generated.

This data set can be created by a user-written program or by using the INSERT function of the MSDB Maintenance utility. Records in the data set are sequenced by MSDB name, and within MSDBs by key.

Related Reading: For a description of the record format and information on how to use the MSDB Maintenance utility, see *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

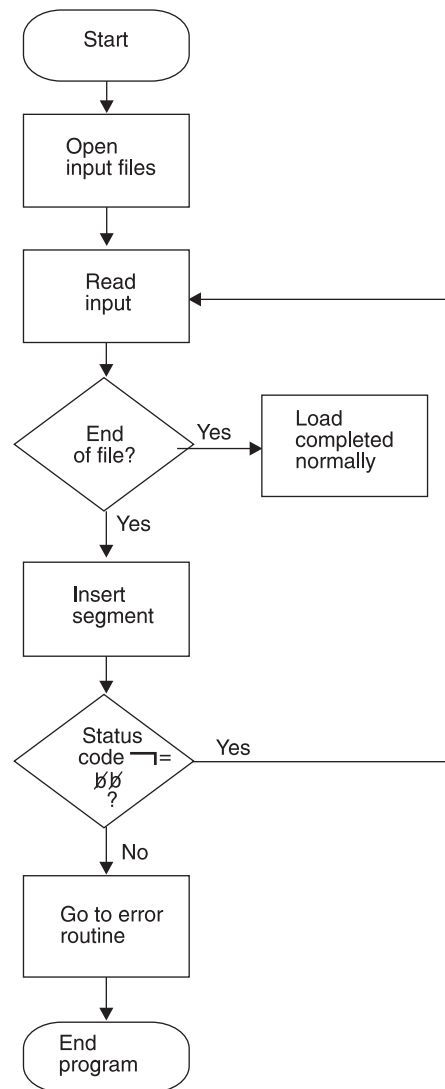


Figure 190. Basic Initial Load Program Logic

Writing a Load Program

```

DLITCBL  START
        PRINT  NOGEN
        SAVE   (14,12),,LOAD1.PROGRAM  SAVE REGISTERS
        USING  DLITCBL,10               DEFINE BASE REGISTER
        LR     10,15                    LOAD BASE REGISTER
        LA     11,SAVEAREA              PERFORM
        ST     13,4(11)                 SAVE
        ST     11,8(13)                  AREA
        LR     13,11                     MAINT
        L      4,0(1)                   LOAD PCB BASE REGISTER
        STCM   4,7,PCBADDR+1            STORE PCB ADDRESS IN CALL LIST
        USING  DLIPCB,4                 DEFINE PCB BASE REGISTER
        OPEN   (LOAD,(INPUT))           OPEN LOAD DATA SOURCE FILE
LOOP     GET    LOAD,CARDAREA            GET SEGMENT TO BE INSERTED
INSERT   CALL   CBLTDLI,MF=(E,DLILINK)  INSERT THE SEGMENT
        AP     SEGCOUNT,=P'1'         INCREMENT SEGMENT COUNT
        CLC    DLISTAT,=CL2' '        WAS COMPLETION NORMAL?
        BE     LOOP                     YES - KEEP GOING
ABEND    ABEND  8,DUMP                   INVALID STATUS
EOF      WTO    'DATABASE 1 LOAD COMPLETED NORMALLY'
        UNPK   COUNTMSG,SEGCOUNT      UNPACK SEGMENT COUNT FOR WTO
        OI     COUNTMSG+4,X'F0'        MAKE SIGN PRINTABLE
        WTO    MF=(E,WTOLIST)          WRITE SEGMENT COUNT
        CLOSE  (LOAD)                  CLOSE INPUT FILE
        L      13,4(13)                 UNCHAIN SAVE AREA
        RETURN (14,12),RC=0            RETURN NORMALLY
        LTORG
SEGCOUNT DC PL3'0'
        DS     0F
WTOLIST   DC AL2(LSTLENGT)
        DC     AL2(0)
COUNTMSG DS CL5
        DC     C' SEGMENTS PROCESSED'
LSTLENGT  EQU (*-WTOLIST)
DLIFUNC   DC CL4'ISRT'                 FUNCTION CODE
DLILINK   DC A(DLIFUNC)                DL/I CALL LIST
PCBADDR   DC A(0)
        DC     A(DATAAREA)
        DC     X'80',AL3(SEGNAME)
CARDAREA  DS 0CL80                     I/O AREA
SEGNAME   DS CL9
SEGKEY    DS 0CL4
DATAAREA  DS CL71
SAVEAREA  DC 18F'0'
LOAD      DCB DDNAME=LOAD1,DSORG=PS,EODAD=EOF,MACRF=(GM),RECFM=FB
DLIPCB    DSECT ,                      DATABASE PCB
DLIDBNAM  DS CL8
DLISGLEV  DS CL2
DLISTAT   DS CL2
DLIPROC   DS CL4
DLIRESV   DS F
DLISEGFB  DS CL8
DLIKEYLN  DS CL4
DLINUMSG  DS CL4
DLIKEYFB  DS CL12
        END

```

Figure 191. Sample Load Program

Restartable Initial Load Program

You should write a restartable initial load program (one that can be restarted from the last checkpoint taken) when the volume of data you need to load is great enough that you would be seriously set back if problems occurred during program

execution. If problems occur and your program is not restartable, the entire load program has to be rerun from the beginning.

Restartable load programs differ from basic load programs in their logic. Figure 192 on page 334 shows the logic for developing a restartable initial load program. If you already have a basic load program, usually only minor changes are required to make it restartable. The basic program must be modified to recognize when restart is taking place, when WTOR requests to stop processing have been made, and when checkpoints have been taken.

Related Reading: Detailed guidance information on what must be done to run a restartable load program under the control of UCF is contained in *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

To make your initial database load program restartable under UCF, consider the following points when you are planning and writing the program:

- If a program is being restarted, the PCB status code will contain a UR prior to the issuance of the first DL/I call. The key feedback area will contain the fully concatenated key of the last segment inserted prior to the last UCF checkpoint taken. (If no checkpoints were taken prior to the failure, this area will contain binary zeros.)
- The UCF does not checkpoint or reposition user files. When restarting, it is the user's responsibility to reposition all such files.
- When restarting, the first DL/I call issued must be an insert of a root segment. For HISAM and HIDAM Index databases, the restart will begin with a GN and a VSAM ERASE sequence to reinsert the higher keys. The resume operation then takes place. Space in the KSDS is reused (recovered) but not in the ESDS. For HDAM, the data will be compared if the root sequence field is unique and a root segment insert is done for a segment that already exists in the database because of segments inserted after the checkpoint. If the segment data is the same, the old segment will be overlaid and the dependent segments will be dropped since they will be reinserted by a subsequent user/reload insert. This occurs only until a non-duplicate root is found. Once a segment with a new key or with different data is encountered, LB status codes will be returned for any subsequent duplicates. Therefore, space is reused for the roots, but lost for the dependent segments.
For HDAM with non-unique keys, any root segments that were inserted after the checkpoint at which the restart was made will remain in the database. This is also true for their dependent segments.
- When the stop request is received, UCF will take a checkpoint just prior to inserting the next root. If the application program fails to terminate, it will be presented the same status code at each of the following root inserts until normal termination of the program.
- For HISAM databases, the RECOVERY option must be specified. For HD organizations, either RECOVERY or SPEED can be defined to Access Method Services.
- UCF checkpoints are taken when the checkpoint count (CKPNT=) has expired and a root insert has been requested. The count refers to the number of root segments inserted and the checkpoint is taken immediately prior to the insertion of the root.

Writing a Load Program

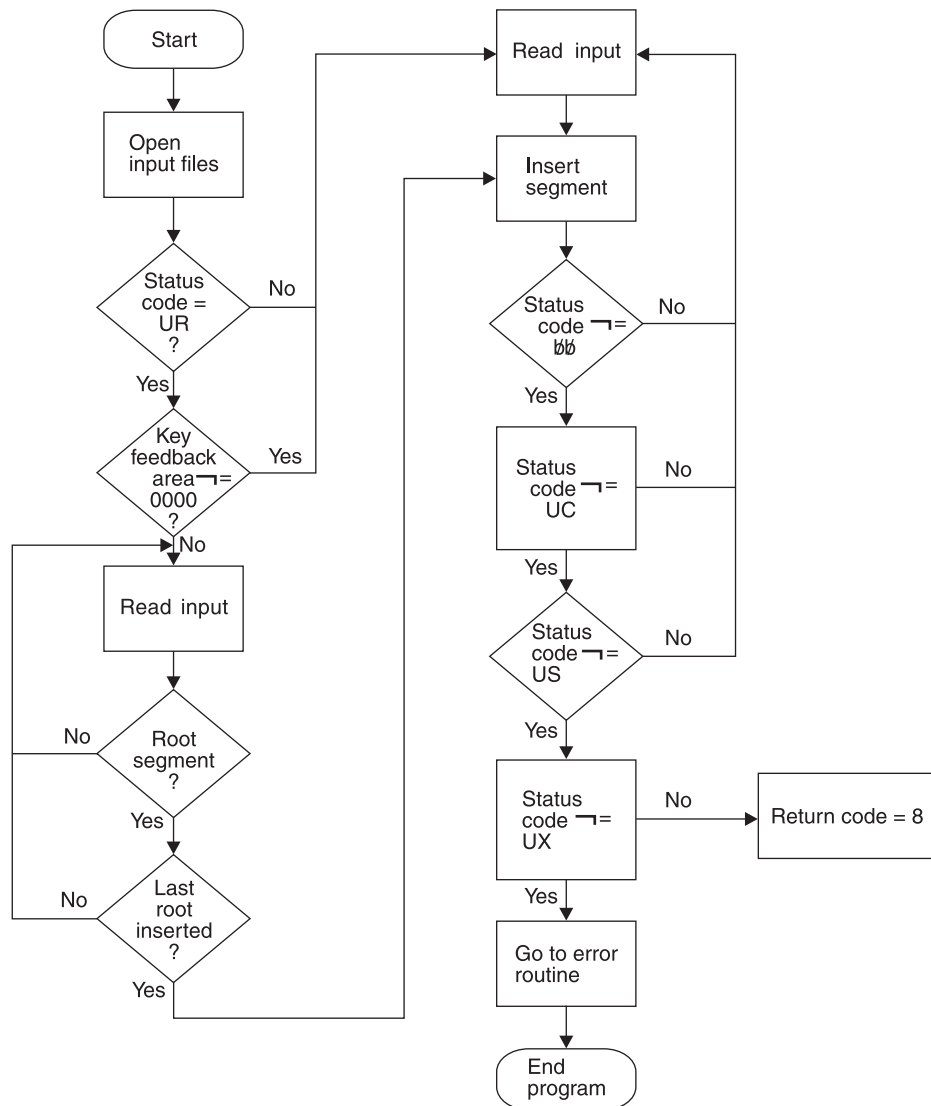


Figure 192. Restartable Initial Load Program Logic

The following lists explains the status codes shown in Figure 192:

- UR** Load program being restarted under control of UCF
- UC** Checkpoint record written to UCF journal data set
- US** Initial load program prepared to stop processing
- UX** Checkpoint record was written and processing stopped

| | | | |
|---------|--------|---|---------------------------------|
| DLITCBL | START | | |
| | PRINT | NOGEN | |
| | SAVE | (14,12),,LOAD1.PROGRAM | SAVE REGISTERS |
| | USING | DLITCBL,10 | DEFINE BASE REGISTER |
| | LR | 10,15 | LOAD BASE REGISTER |
| | LA | 11,SAVEAREA | PERFORM |
| | ST | 13,4(11) | SAVE |
| | ST | 11,8(13) | AREA |
| | LR | 13,11 | MAINT |
| | L | 4,0(1) | LOAD PCB BASE REGISTER |
| | STCM | 4,7,PCBADDR+1 | STORE PCB ADDRESS IN CALL LIST |
| | USING | DLIPCB,4 | DEFINE PCB BASE REGISTER |
| | OPEN | (LOAD,(INPUT)) | OPEN LOAD DATA SOURCE FILE |
| | CLC | DLISTAT,=C'UR' | IS THIS A RESTART? |
| | BNE | NORMAL | NO - BRANCH |
| | CLC | DLIKEYFB(4),=X'00000000' | IS KEY FEEDBACK AREA ZERO? |
| | BE | NORMAL | YES - BRANCH |
| RESTART | WTO | 'RESTART LOAD PROCESSING | FOR DATABASE 1 IS IN PROCESS' |
| RLOOP | GET | LOAD,CARDAREA | GET A LOAD RECORD |
| | CLC | SEGNAME(8),=CL8'SEGMA' | IS THIS A ROOT SEGMENT RECORD? |
| | BNE | RLOOP | NO - KEEP LOOKING |
| | CLC | DLIKEYFB(4),SEGKEY | IS THIS THE LAST ROOT INSERTED? |
| | BNE | RLOOP | NO - KEEP LOOKING |
| | B | INSERT | GO DO IT |
| NORMAL | WTO | 'INITIAL LOAD PROCESSING | FOR DATABASE 1 IS IN PROCESS' |
| LOOP | GET | LOAD,CARDAREA | GET SEGMENT TO BE INSERTED |
| INSERT | CALL | CBLTDLI,MF=(E,DLILINK) | INSERT THE SEGMENT |
| | AP | SEGCOUNT,=P'1' | INCREMENT SEGMENT COUNT |
| | CLC | DLISTAT,=CL2' ' | WAS COMPLETION NORMAL? |
| | BE | LOOP | YES - KEEP GOING |
| | CLC | DLISTAT,=CL2'UC' | HAS CHECKPOINT BEEN TAKEN? |
| | BNE | POINT1 | NO - KEEP CHECKING |
| POINT0 | WTO | 'UCF CHECKPOINT TAKEN FOR | LOAD 1 PROGRAM' |
| | UNPK | COUNTMSG,SEGCOUNT | UNPACK SEGMENT COUNT FOR WTO |
| | OI | COUNTMSG+4,X'F0' | MAKE SIGN PRINTABLE |
| | WTO | MF=(E,WTOLIST) | WRITE SEGMENT COUNT |
| | B | LOOP | NO - KEEP GOING |
| POINT1 | CLC | DLISTAT,=CL2'US' | HAS OPERATOR REQUESTED STOP? |
| | BNE | POINT2 | NO - KEEP CHECKING |
| | B | LOOP | KEEP GOING |
| POINT2 | CLC | DLISTAT,=CL2'UX' | COMBINED CHECKPOINT AND STOP? |
| | BNE | ABEND | NO - GIVE UP |
| | WTO | 'LOAD1 PROGRAM STOPPING PER OPERATOR REQUEST' | |
| | B | RETURN8 | |
| ABEND | ABEND | 8,DUMP | INVALID STATUS |
| EOF | WTO | 'DATABASE 1 LOAD COMPLETED NORMALLY' | |
| | UNPK | COUNTMSG,SEGCOUNT | UNPACK SEGMENT COUNT FOR WTO |
| | OI | COUNTMSG+4,X'F0' | BLAST SIGN |
| | WTO | MF=(E,WTOLIST) | WRITE SEGMENT COUNT |
| | CLOSE | (LOAD) | CLOSE INPUT FILE |
| | L | 13,4(13) | UNCHAIN SAVE AREA |
| | RETURN | (14,12),RC=0 | RETURN NORMALLY |
| RETURN8 | WTO | 'DATABASE 1 LOAD STOPPING | FOR RESTART' |
| | UNPK | COUNTMSG,SEGCOUNT | UNPACK SEGMENT COUNT FOR WTO |
| | OI | COUNTMSG+4,X'F0' | BLAST SIGN |
| | WTO | MF=(E,WTOLIST) | WRITE SEGMENT COUNT |
| | CLOSE | (LOAD) | CLOSE INPUT FILE |
| | L | 13,4(13) | UNCHAIN SAVE AREA |
| | RETURN | (14,12),RC=8 | RETURN AS RESTARTABLE |
| | LTORG | | |

Figure 193. Sample Restartable Initial Load Program (Part 1 of 2)

Writing a Load Program

```
SEGCOUNT DC      PL3'0'  
          DS      0F  
WTOLIST   DC      AL2(LSTLENGT)  
          DC      AL2(0)  
COUNTMSG DS      CL5  
          DC      C' SEGMENTS PROCESSED'  
LSTLENGT EQU      (*-WTOLIST)  
DLIFUNC   DC      CL4'ISRT'          FUNCTION CODE  
DLILINK   DC      A(DLIFUNC)         DL/I CALL LIST  
PCBADDR   DC      A(0)  
          DC      A(DATAAREA)  
          DC      X'80',A13(SEGNAME)  
CARDAREA  DS      0CL80             I/O AREA  
SEGNAME   DS      CL9  
SEGKEY    DS      0CL4  
DATAAREA  DS      CL71  
SAVEAREA  DC      18F'0'  
STOPNDG   DC      X'00'  
LOAD      DCB      DDNAME=LOAD1,DSORG=PS,EODAD=EOF,MACRF=(GM),RECFM=FB  
DLIPCB    DSECT      DATABASE PCB  
DLIDBNAM  DS      CL8  
DLISGLEV  DS      CL2  
DLISTAT   DS      CL2  
DLIPROC   DS      CL4  
DLIRESV   DS      F  
DLISEGFB  DS      CL8  
DLIKEYLN  DS      CL4  
DLINUMSG  DS      CL4  
DLIKEYFB  DS      CL12  
END
```

Figure 193. Sample Restartable Initial Load Program (Part 2 of 2)

JCL for the Initial Load Program

Figure 194 shows the JCL you will need to initially load your database. The //DFSURWF1 DD statement is present only if a logical relationship or secondary index exists.

```
//          EXEC  PGM=DFSRR00,PARM='DLI,your initial load program name,  
//          your PSB name'  
//DFSRESLB  DD    references an authorized library that contains IMS  
              SVC modules  
//STEPLIB   DD    references library that contains your load program  
//          DD    DSN=IMS.SDFSRESL  
//IMS       DD    DSN=IMS.PSBLIB,DISP=SHR  
//          DD    DSN=IMS.DBDLIB,DISP=SHR  
//DFSURWF1  DD    DCB=(RECFM=VB,LRECL=300,  
//          BLKSIZE=(you must specify),  
//          DSN=WF1,DISP=(MOD,PASS)  
//DBNAME    DD    references the database data set to be  
              initially loaded or referenced by  
              the initial load program  
//INPUT     DD    input to your initial load program  
//DFSVSAMP  DD    input for VSAM and OSAM buffers and options  
:  
//*
```

Figure 194. JCL used to initially load a database

Loading a HISAM Database

Segments in a HISAM database are stored in the order in which you present them to the load program. You must present all occurrences of the root segment in

ascending key sequence and all dependent segments of each root in hierarchic sequence. PROCOPT=L (for load) must be specified in the PCB.

Loading a SHISAM Database

Segments in a SHISAM database are stored in the order in which you present them to the load program. You must present all occurrences of the root segment in ascending key sequence. PROCOPT=L (for load) must be specified in the PCB.

Loading a GSAM Database

GSAM databases use logical records, not segments or database records. GSAM logical records are stored in the order in which you present them to the load program.

Loading an HDAM or a PHDAM Database

In an HDAM or a PHDAM database, the user randomizing module determines where a database record is stored, so the sequence in which root segments are presented to the load program does not matter. All dependents of a root should follow the root in hierarchic sequence. PROCOPT=L (for load) or PROCOPT=LS (for load segments in ascending sequence) must be specified in the PCB.

Loading a HIDAM or a PHIDAM Database

To load a HIDAM or a PHIDAM database, you must present root segments in ascending key sequence and all dependents of a root should follow the root in hierarchic sequence. PROCOPT=LS (for load segments in ascending sequence) must be specified in the PCB.

Loading a Database with Logical Relationships or Secondary Indexes

If you are loading a database with logical relationships or secondary indexes, you will need to run, in addition to your load program, some combination of the reorganization utilities. You need to run them to put the correct pointer information in each segment's prefix. These reorganization utilities are described in Chapter 15, "Tuning Databases," on page 347.

Loading Fast Path Databases

This topic describes how to load MSDBs, DEDBs, and sequential dependent segments.

Loading an MSDB

Because MSDBs reside in main storage, you do not load them as you do other IMS databases, that is, by means of a load program that you provide. Rather, they are loaded during system initialization, when they are read from a data set. You first build this data set either by using a program you provide or by running the MSDB Maintenance utility.

Related Reading:

- See *IMS Version 9: Utilities Reference: Database and Transaction Manager* for information on how to the MSDB Maintenance utility.
- See Figure 73 on page 134 for the record format of the MSDBINIT data set.

Loading a DEDB

You load data into a DEDB database with a load program similar to that used for loading other IMS databases. Unlike other load programs, this program runs as a batch message program. The following five steps are necessary to load a DEDB:

1. Calculate space requirements.

The following example assures that root and sequential dependent segment types are loaded in one area.

Assume all root segments are 200 bytes long (198 bytes of data plus 2 bytes for the length field) and that there are 850 root segments in the area. On the average, there are 30 SDEP segments per record. Each is 150 bytes long (148 bytes of data and a 2-byte length field). The CI size is 1024 bytes.

A. Calculate the minimum space required to hold root segments:

| | |
|-------|--|
| 1024 | CI length <i>minus</i> |
| 21 | CI control fields |
| <hr/> | |
| 1003 | equals amount of space for root segments and their prefixes. |

$1003 / 214 = 4.6$ Amount of root and root prefix space divided by length of one root with its prefix equals the number of segments that will fit in one CI.
DEDB segments do not span CIs.
Therefore, only four roots will fit in a CI.

$850 / 4 = 212.5$ The minimum amount of space to hold the defined number of roots to be inserted in this area (850) requires 213 CIs.

After choosing a UOW size, you can determine the DBD specifications for the root addressable and independent overflow parts using the result of the above calculation as a base.

B. Calculate the minimum space required to hold the sequential dependent segments:

| | |
|-------|--|
| 1024 | CI length <i>minus</i> |
| 17 | CI control fields |
| <hr/> | |
| 1007 | equals amount of space for sequential dependents and their prefixes. |

$1007 / 160 = 6.2$ Amount of sequential dependent and prefix space divided by length of one sequential dependent plus its prefix equals the number of segments that will fit in one CI.
Six SDEP segments will fit in a CI.

$30 / 6 = 5$ CIs Minimum amount of space required to hold 30 sequential dependent segments from one root. For 850 roots, the minimum amount of space required is $850 * 5 = 4250$ CIs.

C. Factor into your calculations additional space to take into account:

- The “reorganization UOW”, which is the same size as a regular UOW
- Two control data CIs allocated at the beginning of the root addressable part

- One control data CI for each 120 CIs in the independent overflow part
Assuming a UOW size of 20 CIs, the minimum amount of space to be allocated is: $213 + 4250 + 20 + 2 + 1 = 4486$ CIs.
2. Set up the DBD specifications according to the above results, and execute the DBD generation.
 3. Allocate the VSAM cluster using VSAM Access Method Services.

The following example shows how to allocate an area that would later be referred to as AREA1 in a DBDGEN:

```
DEFINE -
  CLUSTER -
    (NAME (AREA1) -
    VOLUMES (SER123) -
    NONINDEXED -
    CYLINDERS (22) -
    CONTROLINTERVALSIZE (1024) -
    RECORDSIZE (1017) -
    SPEED) -
  DATA -
    (NAME(DATA1)) -
    CATALOG(USERCATLG)
```

The following keywords have special significance when defining an area:

| | |
|----------------------------|---|
| NAME | The name supplied for the cluster is the name subsequently referred to as the area name. The name for the data component is optional. |
| NONINDEXED | DEDB areas are non-indexed clusters. |
| CONTROLINTERVALSIZE | The value supplied, because of a VSAM ICIP requirement, must be 512, 1024, 2048, or 4096. |
| RECORDSIZE | The record size is 7 less than the CI size. These 7 bytes are used for VSAM control information at the end of each CI. |
| SPEED | This keyword is recommended for performance reasons. |
| CATALOG | This optional parameter can be used to specify a user catalog. |

4. Run the DEDB initialization utility (DBFUMINO).
This offline utility must be run to format each area to DBD specifications. Root-addressable and independent-overflow parts are allocated accordingly. The space left in the VSAM cluster is reserved for the sequential-dependent part. Up to 2048 areas can be specified in one utility run; however, the area initializations are serialized. After the run, check the statistical information report against the space calculation results.
5. Run the user DEDB load program.
A BMP program is used to load the DEDB. The randomizing routine used during the loading of the DEDB might have been tailored to direct specific ranges of data to specific areas of the DEDB.
If the load operation fails, the area must be scratched, reallocated, and initialized.

Loading Sequential Dependent Segments

If the order of sequential dependent segments is important, you must consider the way sequential dependents might be loaded in a DEDB. The two alternatives are:

Loading Fast Path Databases

- Add a root and its sequential dependents.

All the sequential dependents of a root are physically written together, but their physical order does not reflect the original data entry sequence. This reflection is not necessarily the way the application needs to view the dependent segments if they are being used primarily as a journal of transactions.

- Add all roots and then the sequential dependents.

This technique restores the SDEP segments to their original entry sequence order. However, it requires a longer process, because the addition of each SDEP segment causes the root to be accessed.

Chapter 14. Monitoring Databases

This chapter describes a number of IMS tools you can use to monitor the performance of your databases. Several tools this chapter does not discuss, but which you can also use for monitoring purposes include:

- IMS Performance Analyzer
- IMS DB Control Suite (On-demand Space Monitor)
- IMS DB Tools Space Monitor Utilities
- DB Integrity Control Facility

Related Reading:

- For information about these and other IMS tools, go to www.ibm.com/ims and link to the IBM DB2 and IMS Tools Web site.
- For information about using the IMS Monitor is found in *IMS Version 9: Administration Guide: System*.
- Additional information about monitoring can also be found in the topic on data sharing in *IMS Version 9: Administration Guide: System*.

In this Chapter:

- “IMS Monitor”
- “Monitoring Fast Path Systems” on page 343

IMS Monitor

The IMS Monitor is a tool that records data about the performance of your DL/I databases in a batch environment. The recorded data is produced in a variety of reports. The monitor’s usefulness is twofold. First, when you run the monitor routinely, it gives you performance data over time. By comparing this data, you can determine whether the performance trend is acceptable. This helps you make decisions about tuning your database and determining when it needs to be reorganized.

The second use of the monitor is to assess how the changes you make effect performance. Once you have accumulated reports describing normal database processing, you can use them as a profile against which to compare the effect of your changes. Examples of changes you might make (then test for performance) include:

- Changes in the structure of your databases
- A change from one DL/I access method to another
- A change in database buffer pool number and size
- Changes in application program logic

In all these cases, your primary goal is probably to minimize the number of I/Os required to perform an operation. The monitor helps you determine whether you have met your objective.

The following example shows how to use the IMS Monitor: suppose you are performing a final test on a new or revised application. The monitor reports show that some DL/I calls in the program, which should have required a single I/O retrieval, actually required a large database scan involving many I/Os. You might be able to correct this problem by making changes in the application program logic.

IMS Monitor

The monitor itself is actually two programs, as shown in Figure 195.

- The IMS Monitor (DFSMNTR0)
- The IMS Monitor Report Print utility (DFSUTR20)

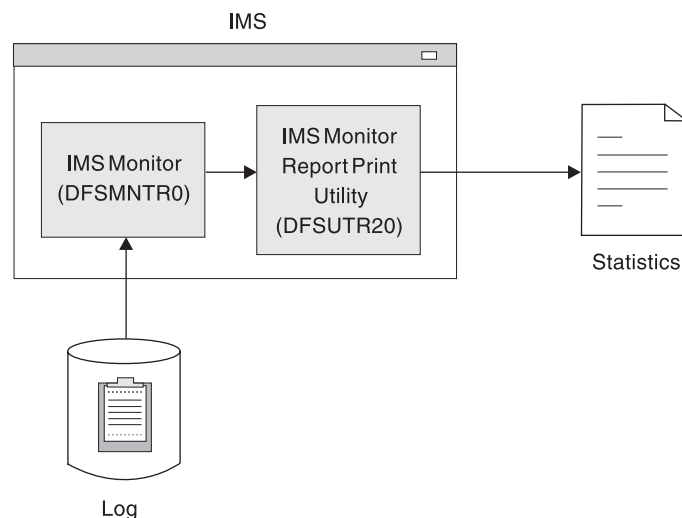


Figure 195. IMS Monitor Works

The IMS Monitor collects data from IMS control blocks (when DL/I is operating) and records the data either on an independent data set or in the IMS log. It collects data with minimum interference to the system. The monitor runs in the same address space as the IMS job, and it can be turned on or off with the MON= parameter in the execution JCL.

The IMS Monitor Report Print utility is an offline program that produces reports summarizing information collected by the IMS Monitor. It produces the following reports:

- VSAM Buffer Pool report
- VSAM Statistics report
- Database Buffer Pool report
- Program I/O report
- DL/I Call Summary report
- Distribution Appendix report
- Monitor Overhead report

Example output of each of these reports is in the *IMS Version 9: Utilities Reference: System*. Each field in the reports is explained, followed by a summary of how you can use the report. Many of these reports are also provided by the IMS Monitor, which is described in *IMS Version 9: Administration Guide: System*. Where the same report is produced by both the DB and IMS Monitor, the description of the report in the *IMS Version 9: Utilities Reference: System* is applicable for both.

Information on operating the IMS Monitor is contained in *IMS Version 9: Operations Guide*.

When the IMS Monitor is on, it remains on until the batch execution ends, requiring some overhead. It cannot be turned on and off from the system console. To

minimize the monitor's impact, use the IMS Monitor in a single-thread test environment rather than multi-thread application environments.

This ensures that the data gathered by the IMS Monitor can be related to a particular program.

Monitoring Fast Path Systems

The major emphasis for monitoring IMS online systems that include message-driven Fast Path applications is the balance between rapid response and high transaction rates. With Fast Path, performance data is made part of the system log information. Because the bulk of the online traffic is expected to be handled by expedited message handling and not be present on the message queues, the Fast Path Log Analysis utility (DBFULTA0) is the prime tool for monitoring Fast Path applications. The IMS Monitor can also be used to monitor Fast Path systems.

Related Reading: For information on using the IMS Monitor for Fast Path systems, see *IMS Version 9: Utilities Reference: System*.

Use the Fast Path Log Analysis utility (DBFULTA0) to prepare statistical reports for Fast Path based on data recorded on the IMS system log. This utility is offline and produces five reports useful for system installation, tuning, and troubleshooting:

- A detailed listing of exception transactions
- A summary of exception detail by transaction code for MPP (message-processing program) regions
- A summary by transaction code for MPP regions
- A summary of IFP, BMP, and CCTL transactions by PSB name or transaction code
- A summary of the log analysis

Do not confuse this utility with the IMS Monitor or the IMS Log Transaction Analysis utility.

Related Reading:

- For more information on CCTL transactions, see the *IMS Version 9: Customization Guide*.
- For more detailed information on the Fast Path Log Analysis utility, see *IMS Version 9: Utilities Reference: System*.

As an administrator in the Fast Path environment, you should perform tasks, like establishing monitoring strategies, performance profiles, and analysis procedures. This topic highlights how to use the Analysis utility to do these tasks, and suggests some Areas where tuning activities might be valuable.

Fast Path Log Analysis Utility

The Fast Path Log Analysis utility gathers statistics of Fast Path exclusive and potential transactions that are passed to Fast Path dependent regions. It reports information for other PSBs (including Fast Path PCBs and the programs that enter the sync point processing) and produces three types of output:

- Formatted summary and detail reports
- A data set of fixed format records for the total traffic of Fast Path transactions extracted from the system logs that form the input to the utility

Monitoring Fast Path Systems

- A data set of records, in the same format, that are selected based on exception conditions (such as those transactions that exceed a certain fixed response time)

The latter data sets can be analyzed in more detail by your installation's programs. They can also be sorted to group critical transactions or events. The details of the record format and meaning of the fields are given in *IMS Version 9: Utilities Reference: System*.

Fast Path Log Reduction

To reduce log volume you can use the LGNR parameter, which is specified during IMS startup. LGNR indicates the maximum number of DEDB buffer alterations that are held before the entire CI is logged.

Related Reading:

- Another way to reduce log volume is to designate the DEDB as nonrecoverable. No changes to the database are logged and no record of database updates is kept in the DBRC RECON data set. See “Non-Recovery Option” on page 118.
- For more information on log reduction and the LGNR parameter, see *IMS Version 9: Utilities Reference: System*.

Fast Path Transaction Timings

For each Fast Path transaction, four time intervals are separately calculated. Figure 196 shows the boundary events and intervals.




Figure 196. Fast Path Transaction Event Timings

The following list describes the four intervals shown in Figure 196:

1. Input queue time: reflects the transaction input queuing within the balancing group to distribute the work.
2. Process time: records the actual elapsed processing time for the individual transaction.
3. Output queue time: shows the effect of sync point in delaying the output message release until after logging.
4. Output message time: shows the line and device availability for receiving the output message. If the transaction originated from a programmable controller, the output time could reflect a delay in dequeue caused by the output not being acknowledged until the next input.

The sum of the first three intervals is termed the *transit time*. This time is slightly different from a response time, because it excludes the line activity for the message, message formatting, and the input edit processing up to the time the message segment leaves the exit routine.

Monitored Events for Fast Path

The control program automatically collects Fast Path event data during system operation. Table 26 shows the information that is made part of the system log records for each Fast Path transaction.

Table 26. Monitor Data for Fast Path Transactions

| Monitored Data | Message-Driven Region | Other Region |
|----------------------------------|-----------------------|--------------|
| Transit and Output Message Times | x | |
| LTERM Name | x | |
| Routing Code | x | |
| Balancing Group Queue Count | x | |
| Number of DEDB Calls | x | x |
| Number of I/O to DEDB | x | x |
| Number of MSDB Calls | x | x |
| Number of CI Contentions | x | x |
| Number of Buffers Allocated | x | x |
| Number of Waits for Buffer | x | x |
| Sync Point Failure Reason Code | x | x |

Selecting Transactions

The analysis utility lets you select transactions to be reported in detail. You give the transaction code and a transit time that each transaction is to exceed, up to a maximum of 65.5 seconds. Several codes can be selected for each utility run. There is also a way to ask for all transactions that exceed the given transit time. In this case, the individual exception specification overrides the general one.

When you do not need to print all such occurrences of the exceptions, you can give a maximum number of detail records to be printed. The default is 1000 individual records, though you can specify up to 9999999 as the maximum number. When you cut off the number of *printed* records, the data set for the exception records contains all transactions that meet the selection criteria.

You can also specify a start time and end time for the transaction reporting interval. The start time corresponds to the earliest transaction that satisfies the clock time (format HH:MM:SS) specified by a utility input control statement. End time is set by the latest transaction that enters the sync point processing before the ending clock time that is specified on an input control statement.

Another selection technique that is available is to select only non-message-driven transactions for reporting. Use this to look at the activity (occurring against MSDBs or DEDBs) caused by calls from IMS programs or BMPs.

Interpreting Fast Path Analysis Reports

The analysis reports show the origin, database activity, and processing events for each transaction code, although most reported items show average and maximum values. The reports produced are:

- Overall summary by transaction
Summarized by transaction code, the transit times and input/output message lengths are given. The database calls and buffer usage are also included.
- Exception detail

Monitoring Fast Path Systems

For those transactions selected, the terminal origin and routing code are given for each individual occurrence of the transaction. The detail also includes the data appearing in the overall summary.

- Summary of exception detail by transaction code

This report is based on the transactions in the exception report. The items reported are the same as for the overall summary.

- Summary of transactions by PSB

All programs that are in non-message-driven regions, MPP regions, and BMP regions that enter the sync point processing are reported. The items reported are the same as the summary of exception detail.

- Recapitulation of the analysis

This is a documentation aid that gives the grand totals of transactions input to the analysis, and the I/O for online utilities.

The combination of the interval covered by the system log input to the utility and the exception criteria you define in the input control statements determines the content of these reports.

Examples of the reports format and the definition of the items reported can be found in *IMS Version 9: Utilities Reference: System*, within the description of the Fast Path Log Analysis utility.

Following are some suggestions for interpreting the reported events:

- Examine the summary reports and investigate the reasons for sync point failure.
- Examine the summary report to see if buffer usage was consistently under the NBA values. Check all negative differences that indicate the need for overflow buffers to see that they were unusual occurrences.
- Compare the database call counts to those of the expected profile. Select those transactions that show unusual patterns for a run to produce a detailed exception report.
- Examine the balancing group queue counts to see if they are conforming with the scheduling algorithm expectations.

Chapter 15. Tuning Databases

Tune your database either to improve performance or to better use database space. This chapter introduces the reorganization utilities, which you can use to tune your database. The chapter also describes the various types of tuning changes you can make with the reorganization utilities, and also when and how to make the changes.

This chapter examines the following aspects of database tuning:

- “Reorganizing the Database”
- “Reorganizing HALDBs” on page 364, including:
 - “HALDB Offline Reorganization” on page 365
 - “HALDB Online Reorganization” on page 370
- “Changing DL/I Access Methods” on page 394
- “Changing the Hierarchic Structure” on page 408
- “Changing Direct-Access Storage Devices” on page 409
- “Tuning OSAM Sequential Buffering” on page 410
- “Adjusting HDAM and PHDAM Options” on page 411
- “Adjusting Buffers” on page 411
- “Adjusting VSAM Options” on page 414
- “Adjusting OSAM Options” on page 416
- “Changing the Amount of Space Allocated” on page 416
- “Changing Operating System Access Methods” on page 417
- “Changing the Number of Data Set Groups” on page 418
- “Tuning Fast Path Systems” on page 421

Keep in mind that when you tune your database, you are often making more than a simple change to it. For example, you might need to reorganize your database and at the same time change operating system access methods. This chapter has procedures to guide you through making each type of change. If you are making more than one change at a time, you should look at the flowchart, Figure 225 on page 419. When used in conjunction with the individual procedures in this chapter, the flowchart guides you in making some types of multiple changes to the database.

Also, some of the tuning changes you make can affect the logic in application programs. You can often use the dictionary to analyze the affect before making changes. In addition, some changes require that you code new DBDs and PSBs. If you initialize your changes in the dictionary, you can then use the dictionary to help create new DBDs and PSBs.

If you are using data sharing, additional information about tuning is in *IMS Version 9: Administration Guide: System*.

Reorganizing the Database

Reorganizing a database means changing how the data in the database is organized to improve performance. In some cases, reorganizing a database might also refer to modifying the database's structure or the structure of the records and segments in the database. Although this chapter focuses on changing how data is organized, you can use many of the reorganization utilities discussed here to make structural changes as well.

Reorganizing the Database

Two database types, DEDB and HALDB, support online reorganization in addition to the offline methods of reorganization discussed here. For more information on the online reorganization of each of these types of databases, see:

- For HALDB, see “HALDB Online Reorganization” on page 370
- For DEDB, search for High-Speed DEDB Direct Reorganization utility (DBFUHDR0) in *IMS Version 9: Utilities Reference: Database and Transaction Manager*

For full-function databases, CI reclaim can also improve performance of DL/I calls to VSAM key sequenced data sets. CI reclaim reduces the number of empty CIs that are read on DL/I calls. For more information, see “VSAM KSDS CI Reclaim For Full-Function Databases” on page 241.

Related Reading: For information about making structural changes to your database, see Chapter 16, “Modifying Databases,” on page 429.

When You Should Reorganize

You should reorganize your database in the following circumstances:

- Database performance has deteriorated. This can happen either because segments in a database record are stored across too many CIs or blocks, or because you are running out of free space in your database.
- There is too much physical I/O to DASD.
- The database structure has changed. For example, you should reorganize a HALDB partition after changing its boundaries or high key.
- The HDAM or PHDAM randomizer has changed.
- The HALDB Partitions Selection exit routine has changed.

The DB Monitor can aid in monitoring a database to help you determine when it is time to reorganize your database. Information about the DB Monitor is found in Chapter 14, “Monitoring Databases,” on page 341.

Reorganizing Databases Offline

You perform three basic steps when reorganizing a database offline (unless you are not making structural changes to the database, in which case, see Chapter 16, “Modifying Databases,” on page 429):

1. Unloading the existing database.
2. Deleting the old database space and defining new database space. (This practice is always good, but it is only necessary if you have multiple extents or volumes, or are using VSAM.) For VSAM, database space refers to the clusters defined to VSAM for database data sets.
3. Reloading the database.

Protecting Your Database During an Offline Reorganization

When you reorganize your database offline, you delete it. Therefore, you should protect it from system or reorganization failure. You can protect your existing database by renaming the space it occupies and then defining new database space. You should take an image copy of your database as soon as it is reloaded and before any application programs are run against it. Taking an image copy provides you with a backup copy of the database and establishes a point of recovery with DBRC in case of system failure. You can create image copies of your

database using the Database Image Copy utility or the Database Image Copy 2 utility, which are described in detail in *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

Offline Reorganization Utilities

IMS utilities can help you reorganize your database. This topic introduces you to these utilities and explains how they work together.

Related Reading: For more information about reorganization utilities, see the *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

Note the following information about the utilities:

- You can use the following reorganization utilities with HALDB:
 - HD Reorganization Unload utility (DFSURGU0)
 - HD Reorganization Reload utility (DFSURGL0)
 - Database Prereorganization utility (DFSURPR0)
 - HALDB Partition Data Set Initialization utility (DFSUPNT0)

For HALDB, both the Database Prereorganization utility and the HALDB Partition Data Set Initialization utility (DFSUPNT0) initialize partitions.

If you are migrating an HDAM or HIDAM database to HALDB, the Prereorganization utility allows some reuse of your existing JCL by disabling full-function database utilities, such as Scan, Prefix Resolution and Prefix Update, in the DFSURCDS data set. After the database is migrated, you can use the HALDB Partition Data Set Initialization utility, which has additional functions such as unconditional specific partition initialization.

- The utilities cannot be used to reorganize HSAM, SHSAM, or GSAM databases. To reorganize these databases, you must write a program to read the old database and then create a new database.
- You are not required to use these reorganization utilities to reorganize your database. You can write your own programs to unload and reload data. You need to write your own programs only if you are making structural changes to your database that cannot be done using these utilities. Information about when these utilities can be used to make structural changes to a database is contained in Chapter 16, “Modifying Databases,” on page 429.
- Several of the reorganization utilities can be used when initially loading a database. They are not used to load the database but to collect and sort the pointer information needed in a segment’s prefix. Therefore, as you read through the utilities you will find some described as “used for initial load or reorganization”.

The reorganization utilities can be classified into three groups, based on the type of reorganization you plan to do:

- Partial reorganization
- Reorganization using UCF
- Reorganization without UCF

Partial Offline Reorganization

If you are reorganizing an HD database, you can reorganize parts of it, rather than the whole database. You would need to reorganize parts, rather than all of it, for two reasons:

- Only parts of it need to be reorganized.

Reorganizing the Database

- By reorganizing only parts of it, you can break the amount of time it takes to do a total reorganization into smaller pieces.

The utilities you use to do a partial reorganization are:

- The Database Surveyor utility, which helps you determine which parts of your database to reorganize
- The Partial Database Reorganization utility, which does the actual reorganization

HALDB partitions do not support partial offline reorganization.

Offline Reorganization Using UCF

Reorganization can be done using a single program, called the Utility Control Facility (UCF), or by using various combinations of utilities. When UCF is used, it acts as a controller, determining which of the various reorganization utilities need to be executed and then getting them executed. Using UCF reduces the number of JCL statements you must create and eliminates the need to sequence the various utilities for execution. It also reduces the number of decisions operations people must make.

Offline Reorganization Without UCF

When you do not use UCF, reorganization of the database is done using a combination of utilities. Which utilities you need to use, and how many, depends on the type of database and whether it uses logical relationships or secondary indexes.

If your database does not use logical relationships or secondary indexes, you simply run the appropriate unload and reload utilities, which are as follows:

- For HISAM databases, the HISAM Reorganization Unload utility and the HISAM Reorganization Reload utility
- For HIDAM index databases (if reorganized separately from the HIDAM database), the HISAM Reorganization Unload utility and the HISAM Reorganization Reload utility
- For SHISAM, HDAM, and HIDAM databases, the HD Reorganization Unload utility and the HD Reorganization Reload utility

If your database does use logical relationships or secondary indexes, you need to run the HD Reorganization Unload and Reload utilities (even if it is a HISAM database). In addition, you must run a variety of other utilities to collect, sort, and restore pointer information from a segment's prefix. Remember, when a database is reorganized, the location of segments changes. If logical relationships or secondary indexes are used, update prefixes to reflect new segment locations. The various utilities involved in updating segment prefixes are:

- Database Prereorganization utility
- Database Scan utility
- Database Prefix Resolution utility
- Database Prefix Update utility

These utilities can also be used to resolve prefix information during initial load of the database.

In the discussion of the utilities in this section, the four unload and reload utilities are discussed first. The four utilities used to resolve prefix information are then discussed. When reading through the utilities for the first time, you need to understand that, if logical relationships or secondary indexes exist (requiring use of the latter four utilities), the sequence in which operations is as follows:

1. Unload
2. Collect more prefix information
3. Reload
4. Collect more prefix information
5. Updated prefixes

You will find, for instance, that the HD Reorganization Reload utility does not just reload the database if a secondary index or logical relationship exists. It reloads the database using one input as a data set containing some of the prefix information that has been collected. It then produces a data set containing more prefix information as output from the reload. When the various utilities do their processing, they use data sets produced by previously executed utilities and produce data sets for use by subsequently executed utilities. When reading through the utilities, watch the input and output data set names, to understand what is happening.

Figure 197 shows you the sequence in which utilities are executed if logical relationships or secondary indexes exist. Figure 198 on page 353 shows the sequence for these utilities when using HALDB partitions.

Reorganizing the Database

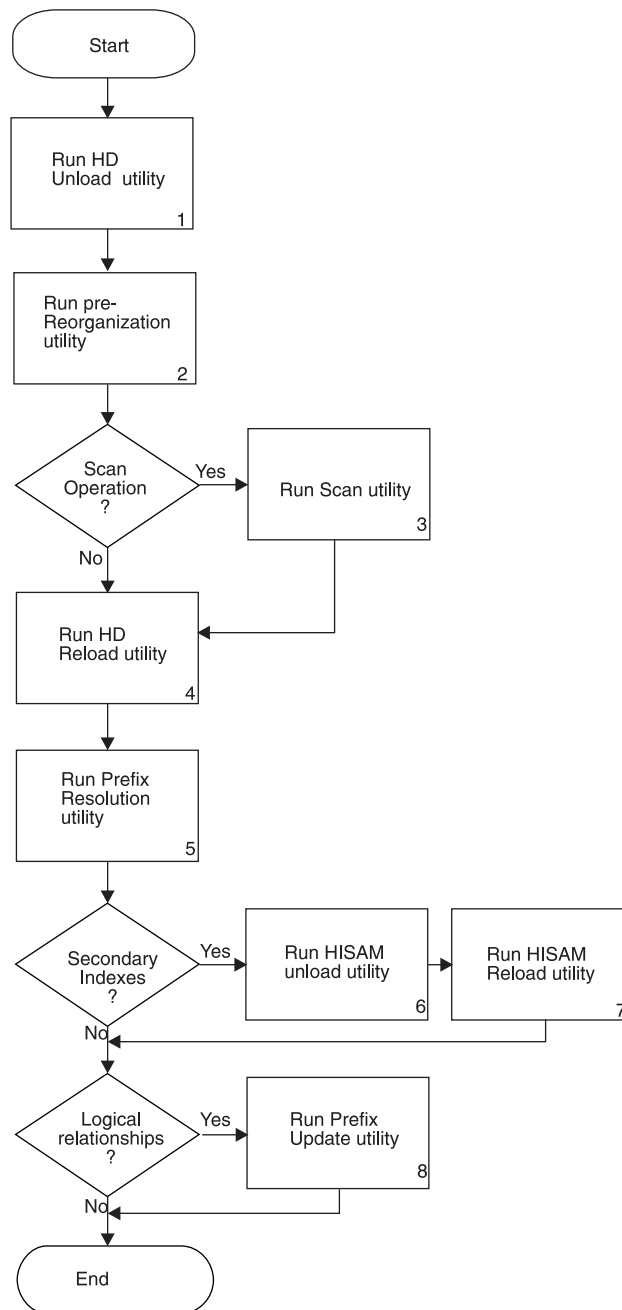


Figure 197. Steps in Reorganizing When Logical Relationships or Secondary Indexes Exist

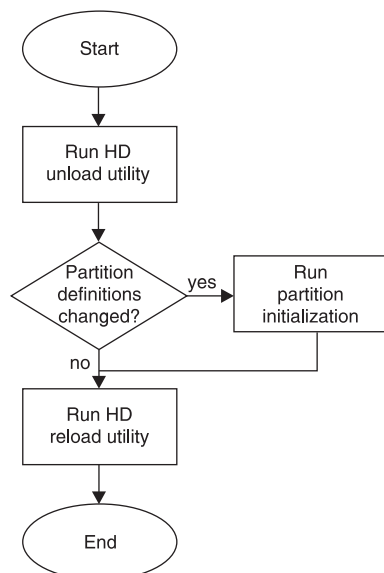


Figure 198. Steps for Reorganizing HALDB Partitions When Logical Relationships or Secondary Indexes Exist

As an alternative, where Figure 198 calls for the Partition Initialization utility, you can run the Prereorganization utility.

HISAM Reorganization Unload Utility (DFSURUL0)

Figure 199 shows the input to and output from the HISAM Reorganization Unload utility.

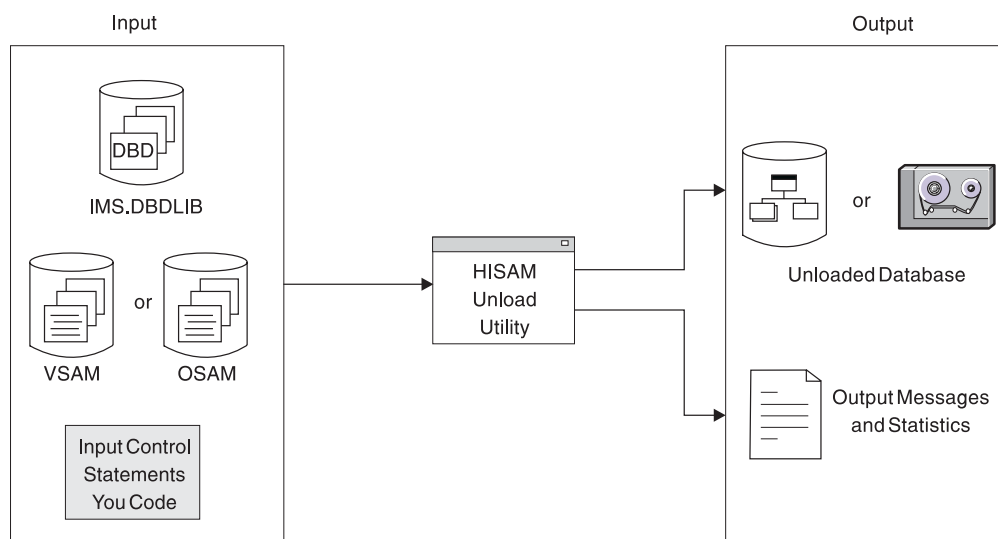


Figure 199. HISAM Reorganization Unload Utility (DFSURUL0)

You use the HISAM Unload utility to unload a HISAM database or HIDAM index database. (SHISAM databases are unloaded using the HD Reorganization Unload utility.) If your database uses secondary indexes, you also use the HISAM Unload utility (later in the reorganization process) to perform a variety of other operations associated with secondary indexes.

HISAM Reorganization Reload Utility (DFSURRL0)

Figure 200 shows the input to and output from the HISAM Reorganization Reload utility.

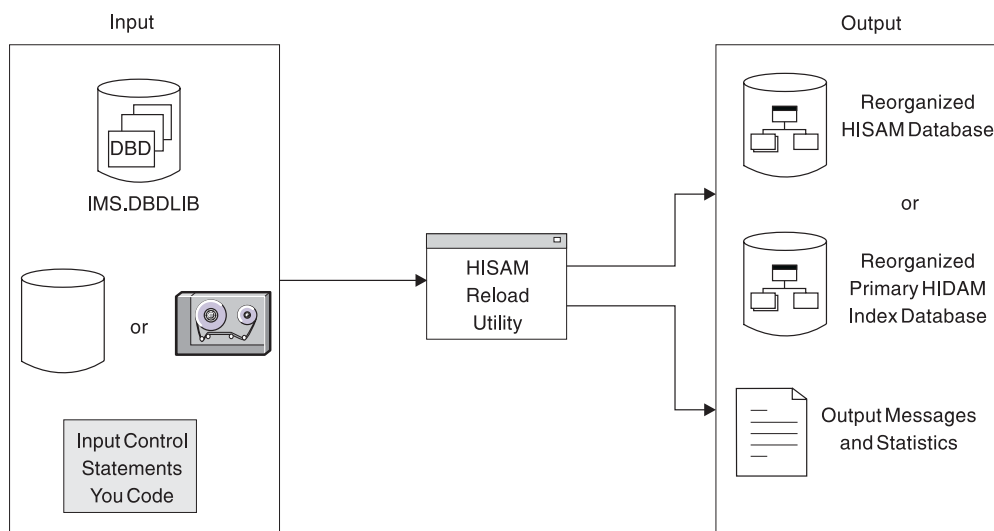


Figure 200. HISAM Reorganization Reload Utility (DFSURRL0)

You use the HISAM reload utility to reload a HISAM database. (SHISAM databases are reloaded using the HD Reorganization Reload utility.) You also use the HISAM reload utility to reload the primary index of a HIDAM database. If your databases use secondary indexes, you use the HISAM reload utility (later in the reorganization process) to perform a variety of other operations associated with secondary indexes.

HD Reorganization Unload Utility (DFSURGU0)

Figure 201 shows the input to and output from the HD Reorganization Unload utility.

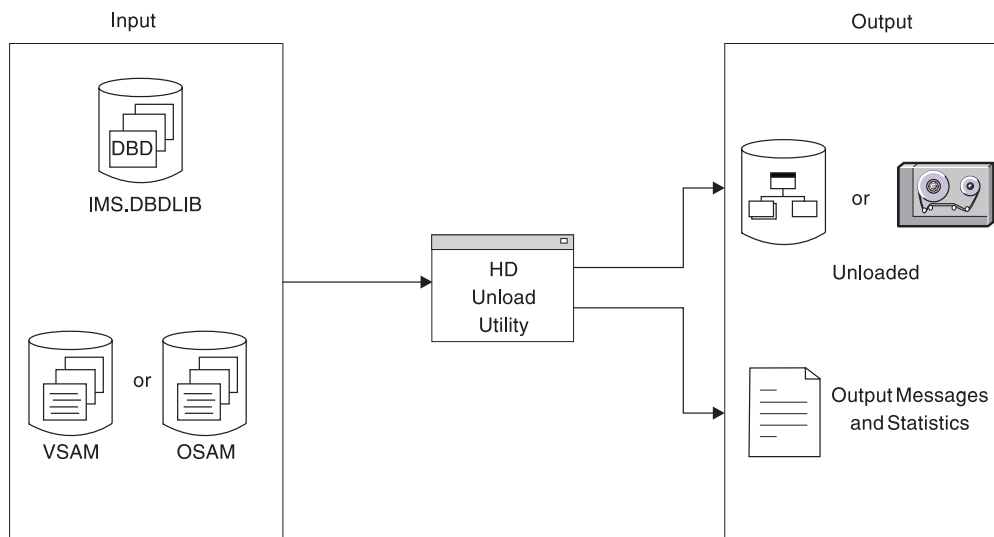


Figure 201. HD Reorganization Unload Utility (DFSURGU0)

You use the HD Reorganization Unload utility to unload:

- HDAM, HIDAM, or SHISAM databases

- HISAM databases that use secondary indexes
- HISAM databases that use symbolic pointers in a logical relationship
- HISAM databases without segment/edit compression that are being converted to HISAM databases with segment/edit compression.
- PHDAM databases or partitions
- PHIDAM databases or partitions
- PSINDEX databases or partitions

If you use the HD Reorganization Unload utility to unload a HALDB (a PHDAM, PHIDAM, or PSINDEX database), you do not need to include DD statements for the database data sets. The HD Reorganization Unload utility uses dynamic allocation for HALDB data sets.

HD Reorganization Reload Utility (DFSURGL0)

Figure 202 shows the input to and output from the HD Reorganization Reload utility.

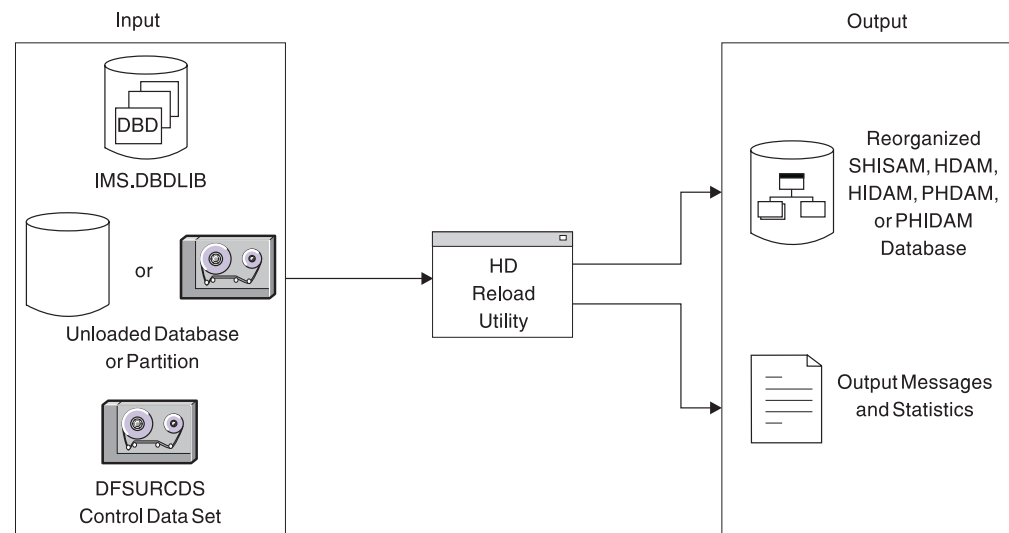


Figure 202. HD Reorganization Reload Utility (DFSURGL0)

You use the HD Reorganization Reload utility to reload:

- HDAM, HIDAM, PHDAM, PHIDAM, PSINDEX, or SHISAM databases
- HISAM databases that use logical relationships or secondary indexes
- HISAM databases without segment/edit compression that are being converted to HISAM databases with segment/edit compression

If logical relationships or secondary indexes exist in the database being reloaded, the DFSURCDS control data set created by the Prereorganization utility is used as one input to the HD Reorganization Reload utility. The DFSURCDS control data set contains information needed to resolve secondary index or logical relationship pointers.

When logical relationships or secondary indexes exist, the HD Reorganization Reload utility produces as output the DFSURWF1 work data set. DFSURCDS identifies the information that will be collected on DFSURWF1.

The DFSURWF1 work data set will become input to the Database Prefix Resolution utility. Note in Figure 202 that, if the database being reloaded has a primary index, it

Reorganizing the Database

is reloaded automatically when the main database is reloaded. A HIDAM index database can also be reorganized as a separate operation using the HISAM unload and reload utilities.

Exception: DFSURWF1 is not used for HALDBs.

Database Prereorganization Utility (DFSURPR0)

Figure 203 shows the input to and output from the Database Prereorganization utility.

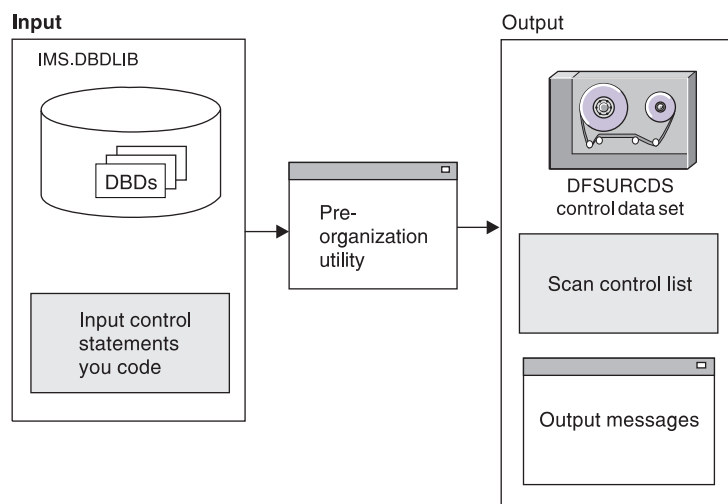


Figure 203. Database Prereorganization Utility (DFSURPR0)

You use the Database Prereorganization utility when:

- A database to be initially loaded or reorganized has secondary indexes or logical relationships
- A database *not* being initially loaded or reorganized contains segments involved in logical relationships with databases that are being loaded or reorganized

The Database Prereorganization utility produces the DFSURCDS control data set, which contains information about what pointers need to be resolved later if secondary indexing or logical relationships exist. The DFSURCDS control data set produced by the Prereorganization utility is used as input to the following:

- The Database Scan utility, if that utility needs to be run
- The HD Reorganization Reload utility, if secondary indexing or logical relationships exist
- The Database Prefix Resolution utility, after the database is loaded or reloaded

The Prereorganization utility also produces a list of which databases not being initially loaded or reorganized contain segments involved in logical relationships with the database that is being initially loaded or reorganized.

This utility is always run *before* the database is loaded (for initial load) or reloaded (for reorganization).

Database Scan Utility (DFSURGS0)

Figure 204 shows the input to and output from the Database Scan utility.

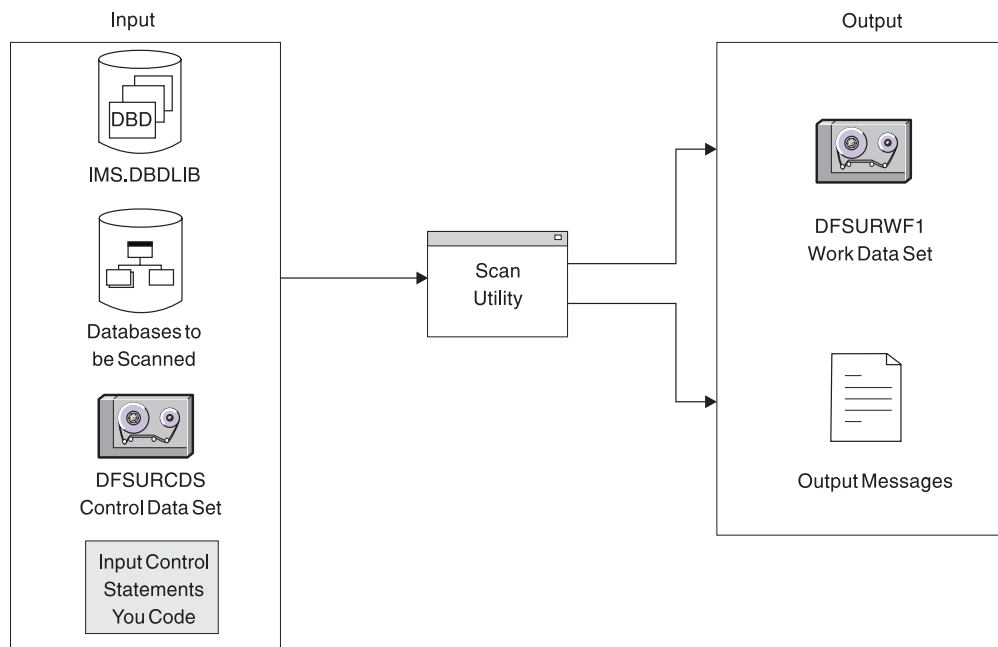


Figure 204. Database Scan Utility (DFSURGS0)

You use the Database Scan utility to scan databases that are not being initially loaded or reorganized but contain segments involved in logical relationships with databases that are being initially loaded or reorganized. For input, the utility uses the DFSURCDS control data set created by the Prereorganization utility. For output, the utility produces the DFSURWF1 work data set, which contains prefix information needed to resolve logical relationships. The DFSURWF1 work data set is used as input to the Database Prefix Resolution utility.

This utility is always run *before* the database is loaded (for initial load) or reloaded (for reorganization).

Database Prefix Resolution Utility (DFSURG10)

Figure 205 shows the input to and output from the Database Prefix Resolution utility.

Reorganizing the Database

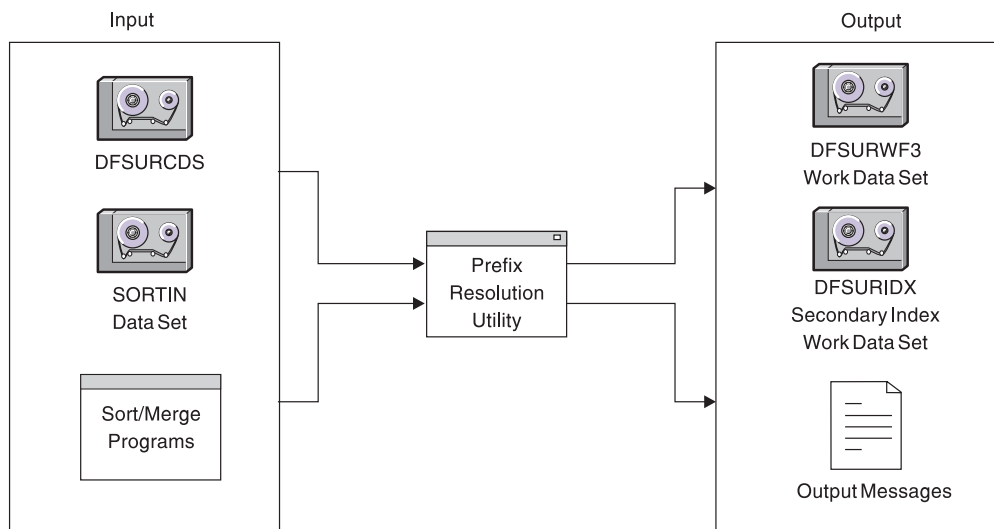


Figure 205. Database Prefix Resolution Utility (DFSURG10)

You use the Prefix Resolution utility to accumulate and sort the information that has been put on DFSURWF1 work data sets up to this point in the load or reload process. The various work data sets that could be input to this utility are:

- The DFSURCDS control data set produced by the Prereorganization utility
- The DFSURWF1 work data set produced by the scan utility
- The DFSURWF1 work data set produced by the HD Reorganization Reload utility

The DFSURWF1 work data sets must be concatenated to form an input data set for the Prefix Resolution utility. The name of the input data set is SORTIN.

The Prefix Resolution utility uses the z/OS sort/merge programs to sort the information that has been accumulated. For output, the utility produces the DFSURWF3 work data set, which contains the sorted prefix information needed to resolve logical relationships. The DFSURWF3 data set will become input to the Database Prefix Update utility.

If secondary indexes exist, the utility produces the DFSURIDX work data set, which contains the information needed to create a new secondary index or update a shared secondary index database. The DFSURIDX work data set is used as input to the HISAM unload utility. The HISAM unload utility formats the secondary index information before the HISAM reload utility creates a secondary index or updates a shared secondary index database.

This utility is always run *after* the database is loaded (for initial load) or reloaded (for reorganization).

Database Prefix Update Utility (DFSURGP0)

Figure 206 shows the input to and output from the Database Prefix Update utility.




Figure 206. Database Prefix Update Utility (DFSURGP0)

You use the Prefix Update utility to update the prefix of each segment whose prefix was affected by the initial loading or reorganization of the database. The prefix fields that are updated include the logical parent, logical twin, and logical child pointer fields, and the counter fields for logical parents. The Prefix Update utility uses as input the DFSURWF3 data set created by the Prefix Resolution utility.

This utility is always run *after* the database is loaded (for initial load) or reloaded (for reorganization) and after the Prefix Resolution utility has been run.

Using HISAM Unload and Reload Utilities for Secondary Indexing Operations

In addition to using the HISAM unload and reload utilities to unload and reload a database, you can also use them to:

- Build a secondary index database
- Merge a secondary index into a shared secondary index database
- Replace a secondary index in a shared secondary index database
- Extract a secondary index from a shared secondary index database

Each of these operations is done separately. That is, none of them can be done in conjunction with running the HISAM unload and reload utilities to unload or reload a regular database.

Figure 207 on page 360 shows the input to and output from the HISAM unload and reload utilities when performing the first three operations. The DFSURIDX work data set used as input to the HISAM unload utility was created by the Prefix Resolution utility. It contains the information needed to create or update a shared secondary index database. The HISAM unload utility formats the secondary index information for use by the HISAM reload utility. Note that the input control statement to the HISAM unload utility has an X in position 1 when the utility is used for secondary indexing operations rather than for unloading a regular database. Position 3 contains one of the following characters:

- M: means the operation is either to build a new secondary index database or merge a secondary index into a shared secondary index database
- R: means the operation is to replace a secondary index into a shared secondary index database

Reorganizing the Database

The HISAM reload utility uses the output from the HISAM unload utility to create the new secondary index or merge or replace the secondary index in a shared secondary index database.

Figure 208 on page 361 shows the input to and output from the HISAM unload utility when an index is being extracted from a set of shared indexes. Note that the input can be one of the following:

- The DFSURIDX work data set created by the Prefix Resolution utility
- The shared secondary index database

Again, position 1 in the input control statement contains an X. Position 3 contains an E, which means the operation is to extract a secondary index.

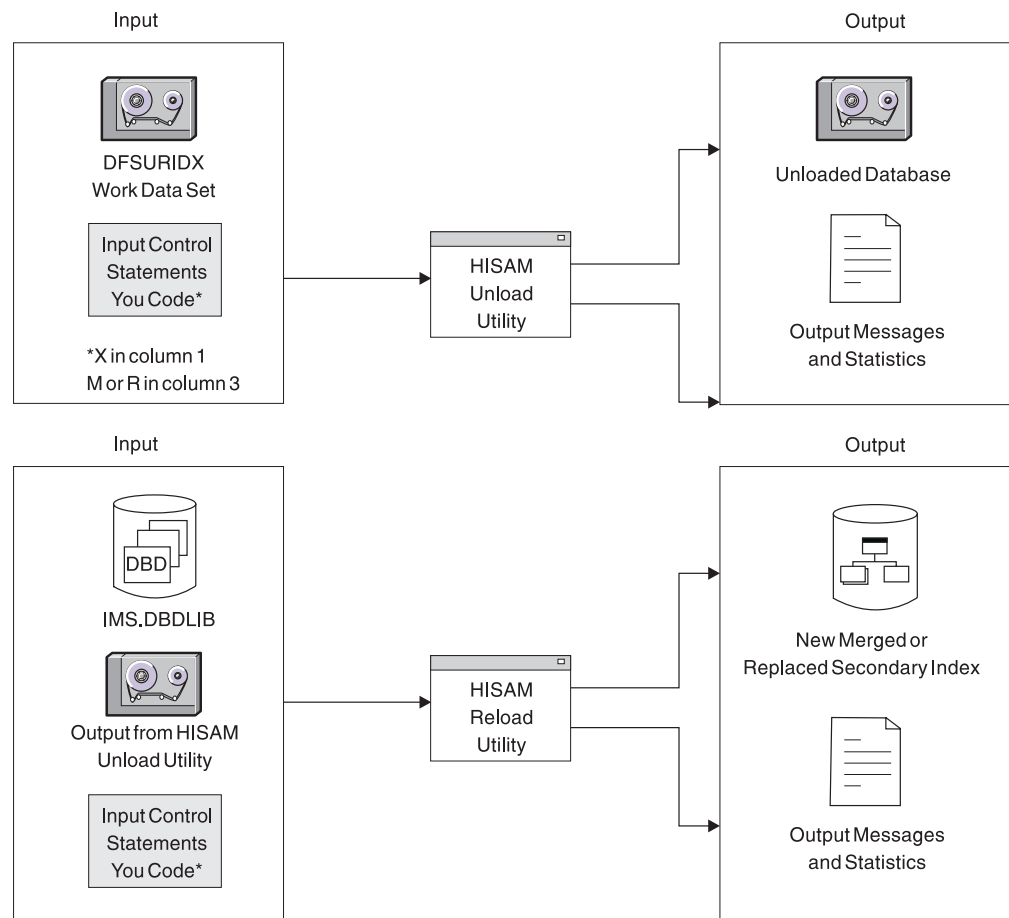


Figure 207. HISAM Reorganization Unload and Reload Utilities Used for Create, Merge, or Replace Secondary Indexing Operations

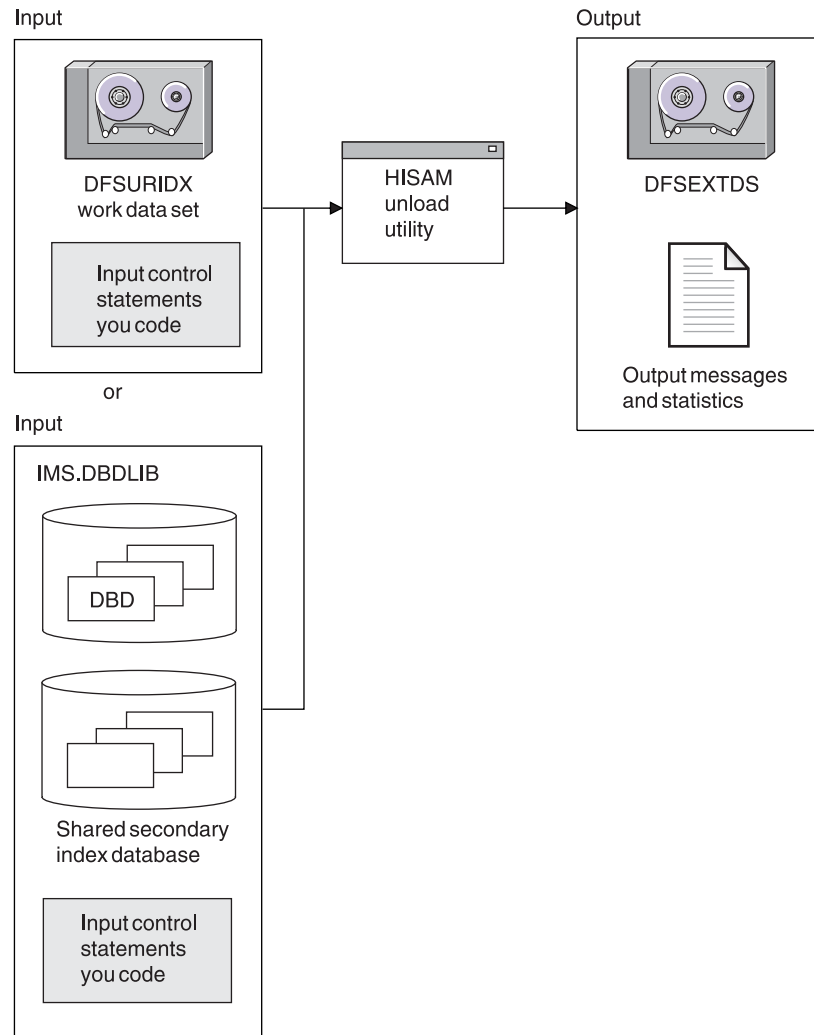


Figure 208. HISAM Reorganization Unload Utility Used for Extract Secondary Indexing Operations

Utility Control Facility (DFSUCF00)

The Utility Control Facility is a program that controls the execution of reorganization and recovery utilities. Control here means that it generates many of the JCL statements you must create and eliminates the need to sequence the various utilities for execution. The only reorganization utilities that cannot be run under the control of UCF are the Database Surveyor utility and the Partial Database Reorganization utility. In addition to controlling the execution of other utilities, UCF allows you to stop and then later restart a job.

Database Surveyor Utility (DFSPRSUR)

Figure 209 on page 362 shows the input to and output from the Database Surveyor utility.

Reorganizing the Database

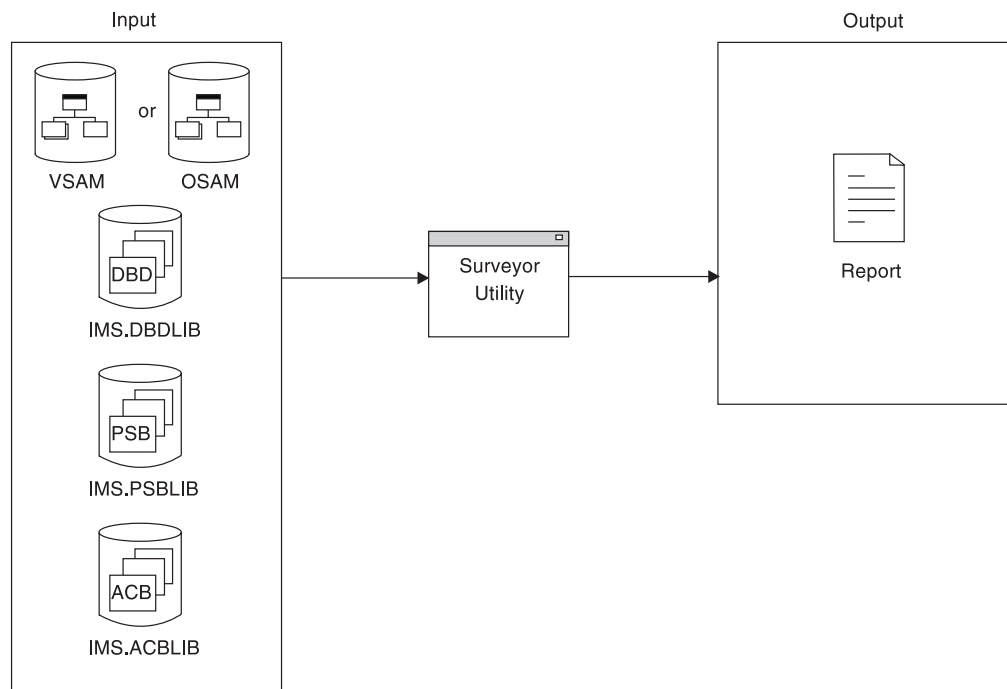


Figure 209. Database Surveyor Utility (DFSPRSUR)

Use the Surveyor utility to scan all or part of an HDAM or a HIDAM database to determine whether a reorganization is needed. The Surveyor utility produces a report describing the physical organization of the database. The report includes the size and location of areas of free space. When you do a partial reorganization, you will know where free space exists into which you can put your reorganized database records.

Partial Database Reorganization Utility (DFSPRCT1)

Figure 210 on page 363 shows the input to and output from the Partial Database Reorganization utility.

You would use the Partial Database Reorganization utility to reorganize parts of your HD database. It can be used when HD databases use secondary indexes or logical relationships. You tell the utility what range of records you need reorganized.

- In an HDAM database, a range is a group of database records with continuous relative block numbers.
- In a HIDAM database, a range is a group of database records with continuous key values.

Generally, before using the Partial Database Reorganization utility, you would run the Database Surveyor utility (described in “Database Surveyor Utility (DFSPRSUR)” on page 361). The Surveyor utility helps you determine whether a reorganization is needed and find the location and size of areas of free space. You need to know the location and size of areas of free space so you will know where to put reorganized database records.

The Partial Database Reorganization utility reorganizes the database in two steps:

1. In the first step, the utility produces control tables for use in Step 2, which is when the actual reorganization is done. As an option, the utility can produce PSB source statements for creating a PSB for use in Step 2. The utility also generates reports that show which logically related segments in logically related

databases must be scanned in Step 2, and which can be optionally scanned in Step 2. (Some GSAM databases are involved in Step 2 for which a PSB is needed.)

2. In the second step, the utility does the actual reorganization. The database records you have specified are unloaded to a data set. The space they occupied in the database is freed. Then database records are reloaded into the database in the range of free space you specified. Finally, all pointers to database records with new locations are changed to point to the new location. A report is produced at the end of Step 2 to tell you what was done.

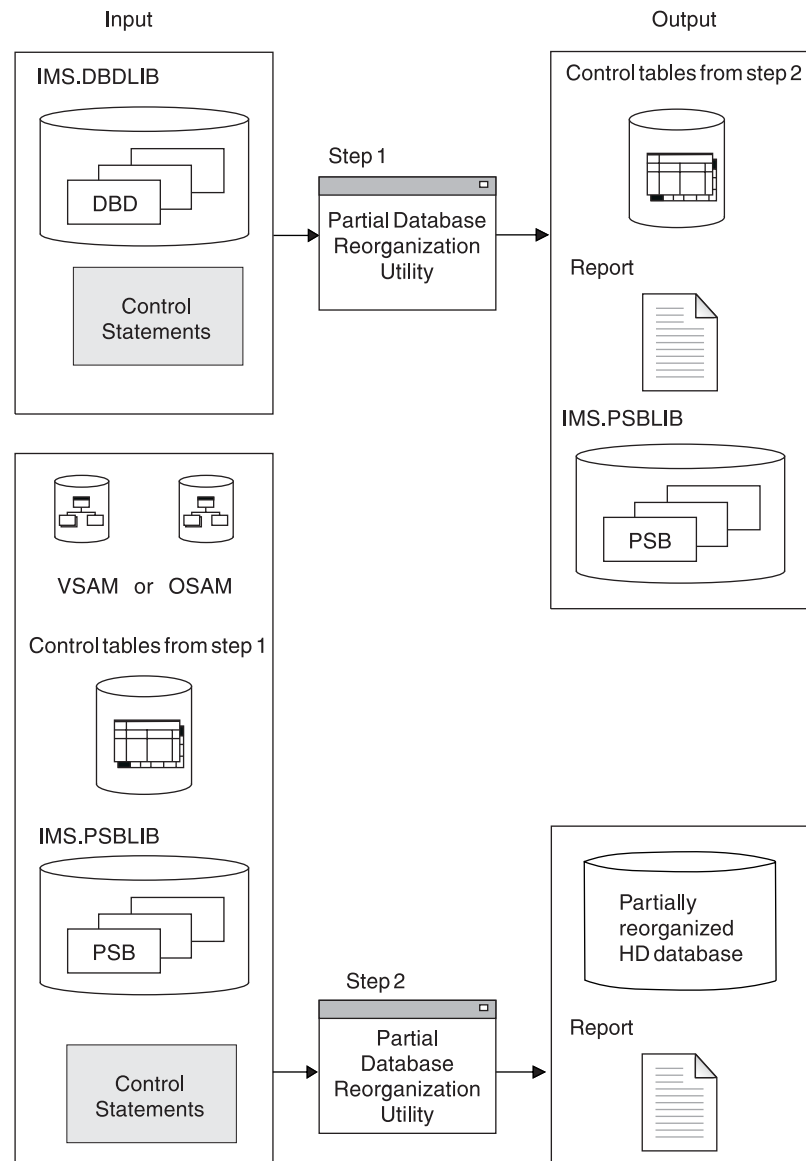


Figure 210. Partial Database Reorganization Utility (DFSPRCT1)

Procedures for Offline Database Reorganizations

This topic describes how to reorganize offline the following database and index types:

- HISAM
- HD (HDAM or HIDAM)

Reorganizing the Database

- Primary or Secondary Index

Reorganizing a HISAM Database (No Secondary Indexes)

To reorganize a HISAM database when it does not use logical relationships or secondary indexes:

1. Unload the database using the HISAM Reorganization Unload utility.
2. Any time you unload a data set, you should delete and reallocate the data set before reloading.
3. Reload the database using the HISAM Reorganization Reload utility. Make an image copy of your database once it is reloaded.

Reorganizing an HD (HDAM or HIDAM) Database (No Logical Relationships or Secondary Indexes)

To reorganize an HD database when it does not use logical relationships or secondary indexes:

1. Unload the database using the HD Reorganization Unload utility.
2. Any time you unload a data set, you should delete and reallocate the data set before reloading.
3. Reload the database using the HD Reorganization Reload utility. Make an image copy of your database once it is reloaded.

Reorganizing a Primary or Secondary Index

HIDAM has a primary index. HISAM, HDAM, and HIDAM have separate secondary index databases when secondary indexing is being used. Both index types are reorganized in the same way:

1. Unload the index database using the HISAM Reorganization Unload utility.
2. Any time you unload a data set, you should delete and reallocate the data set before reloading.
3. Reload the index database using the HISAM Reorganization Reload utility. Make an image copy of your database as soon as it is reloaded.

Reorganizing HALDBs

One of the primary advantages of HALDB is its simplified and shortened reorganization process and the ability to reorganize HALDB databases online using the integrated HALDB Online Reorganization function.

Both PHDAM and PHIDAM HALDBs can be reorganized online or offline. A PSINDEX HALDB can be reorganized only offline. Whether you are reorganizing your HALDB online or offline, the reorganization process is different from the reorganization processes used for other full-function databases.

Reorganizations of HALDBs with logical relationships and secondary indexes do not require the execution of utilities to update these pointers. Instead, HALDB uses a self-healing pointer process to correct these pointers when they are used.

These subjects are discussed in the following topics:

- “HALDB Offline Reorganization” on page 365
- “HALDB Online Reorganization” on page 370
- “The HALDB Self-Healing Pointer Process” on page 388

HALDB Offline Reorganization

The offline reorganization processes for a HALDB database and other full-function IMS databases are similar: they both consist of an unload and reload of the database. The HALDB offline reorganization process has advantages over the reorganization process of other full-function databases, such as:

- You can reorganize one HALDB partition at a time or reorganize multiple partitions in parallel.
- The self-healing pointer process of HALDBs eliminates the need to manually update logical relationships and secondary indexes after reorganizing a HALDB.
- You do not need to include DD statements for HALDB data sets when you reorganize a HALDB. HALDB data sets are dynamically allocated.

Overview of HALDB Offline Reorganization

A offline reorganization of a HALDB database can be done with one or more parallel processes. These processes unload one or more partitions and reload them. If the database has secondary indexes or logical relationships, additional steps are not required. The HALDB self-healing process makes updates of pointers during the reorganization unnecessary. The amount of time required for a reorganization depends on the sizes of the partitions. Smaller partitions reduce the time. You can reduce your reorganization time by creating more partitions and reorganizing them in parallel.

The basic steps involved in reorganizing a HALDB offline are:

1. Run the HD Reorganization Unload utility (DFSURGU0) to unload the entire database, a range of partitions, or a single partition.
2. Optionally, initialize the partitions by running either of the following utilities:
 - HALDB Partition Data Set Initialization utility (DFSUPNT0)
 - Database Prereorganization utility (DFSURPR0)
3. Run the HD Reorganization Reload utility to reload the database or partitions.
4. Make image copies of all reloaded partition data sets.

Figure 211 on page 366 shows the offline processes used to reorganize a HALDB database with logical relationships and secondary indexes. In this case, the partitions are reorganized by parallel processes. Each partition can be unloaded and reloaded in less time than unloading and reloading the entire database. This is much faster than the process for a non-HALDB full-function database. Additionally, no time is required for updating pointers in the logically related database or rebuilding secondary indexes. This further shortens the process.

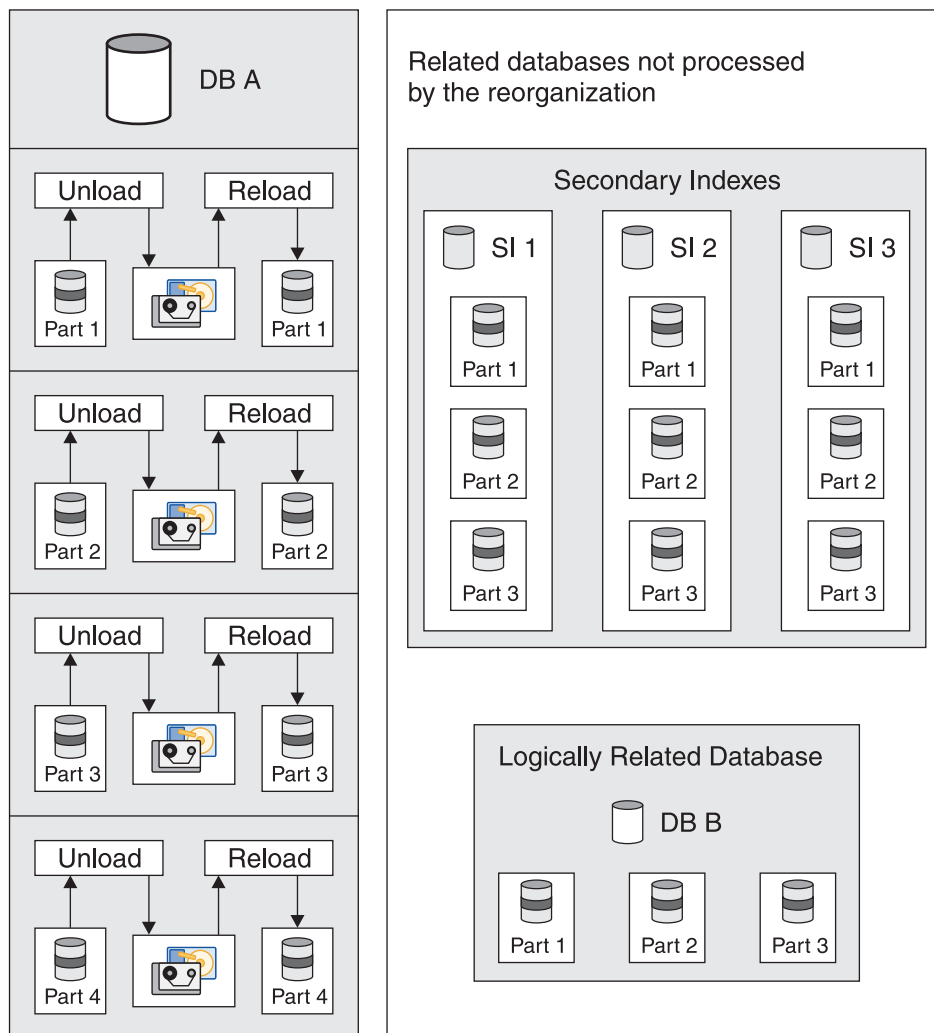


Figure 211. Offline Reorganization of a HALDB database

Related Reading: To compare the HALDB reorganization process illustrated in Figure 211 with the reorganization process for other full-function databases, see the flow chart of the steps for reorganizing non-HALDB databases that use logical relationships or secondary indexes in Figure 197 on page 352.

Options for Offline Reorganization of HALDBs

You have several options when reorganizing a HALDB database:

- You can reorganize any number of partitions. If you need to reorganize only one partition, you can unload and reload it without processing other partitions.
- You can reorganize partitions in parallel or you can reorganize the database with one process. The degree of parallelism is determined by the number of reorganization jobs you run. Each job can process one or multiple partitions. To increase the parallelism, you can increase the number of reorganization jobs and decrease the number of partitions each job processes.
- You can reuse existing database data sets or you can delete them after they are unloaded and allocate new data sets for the reload.
- You can add partitions, delete partitions, or change partition boundaries.

Related Reading:

- For information on changing partition definitions, including partition boundaries, see “Changing HALDB Partition Definitions” on page 404.
- For information on using the Partition Definition utility to change partition definitions in DBRC, see “Creating HALDBs with the HALDB Partition Definition Utility” on page 298.

Unloading HALDB Partitions and Databases for Offline Reorganization

HALDB partitions or databases can be unloaded with the HD Reorganization Unload utility (DFSURGU0).

To unload an entire HALDB database, do not include a SYSIN DD statement. To unload any other number of partitions, you must include a control statement in your SYSIN data set. The control statement identifies the name of the first partition to unload and, if you are unloading more than one partition, the number of partitions to unload.

Multiple partitions are unloaded consecutively. For key range partitioning, consecutive partitions are determined by the high keys. If you are using a partition selection exit routine, consecutive partitions are determined by the order assigned by the exit routine.

Do not include DD statements for the HALDB database data sets. The HD Reorganization Unload utility uses dynamic allocation for HALDB data sets. This is not true for non-HALDB databases.

Requirement: You must supply buffer pools for all data sets in the partitions that are unloaded. This includes the ILDSs.

Recommendation: Enable OSAM sequential buffering for databases that use OSAM.

Figure 212 shows a control statement to unload one partition.

```
SYSIN DD *
PARTITION=PE002
```

Figure 212. Example: The HD Reorganization Unload Utility Control Statement to Unload One Partition

Figure 213 shows a control statement to unload three partitions.

```
SYSIN DD *
PARTITION=PE004,NUMBER=3
```

Figure 213. Example: The HD Reorganization Unload Utility Control Statement to Unload Multiple Partitions

Figure 214 on page 368 shows a sample job that unloads a HALDB partition.

Reorganizing HALDBs

```
//JOUK03C JOB (999,POK),JOUK03,CLASS=A,NOTIFY=&SYSUID,
//      MSGLEVEL=(1,1),MSGCLASS=X,REGION=0M
//JOB LIB DD DSN=IMSPSA.IMS0.SDFSRESL,DISP=SHR
//      DD DSN=IMSPSA.IM0A.MDALIB,DISP=SHR
//*****
//* HD UNLOAD FOR THE PARTITION PE002 OF PEOPLE DATABASE
//*****
//UNLOAD EXEC PGM=DFSRR00,REGION=1024K,
//      PARM='ULU,DFSURGU0,PEOPLE,,,,,,Y,N'
//DFSRESLB DD DSN=IMSPSA.IMS0.SDFSRESL,DISP=SHR
//IMS DD DISP=SHR,DSN=JOUK03.HALDB.DBDLIB
//DFSURGU1 DD DSN=JOUK03.UNLOAD.PE002,UNIT=3390,VOL=SER=TOTIMN,
//      SPACE=(CYL,(10,5),RLSE),DISP=(NEW,CATLG)
//DFSVSAMP DD *
IOBF=(4096,50)
VSRBF=8192,50
/*
//SYSPRINT DD SYSOUT=*
//DFSCCTL DD *
SBPARM ACTIV=COND
/*
//SYSIN DD *
PARTITION=PE002
/*
```

Figure 214. Example: Sample JCL to Unload a HALDB Partition

The IMS High Performance Unload tool is an alternative to the HD Reorganization Unload utility. You can use the High Performance Unload tool to unload any number of partitions or the entire database.

Related Reading: For more information on the High Performance Unload tool, see *IBM DB2 and IMS Tools: IMS High Performance Unload for OS/390*.

Reallocating HALDB Database Data Sets for Offline Reorganization

You do not have to delete and redefine HALDB database data sets before you reload them. This applies to both OSAM and VSAM data sets. VSAM data sets, other than ILDSs, must be specified with the REUSE option. HALDB supports this option.

If you delete and redefine partition data sets, but do not reload data into them, you must initialize the partition data sets. If you reload data into the partition data sets after deleting and redefining them, you do not need to initialize the partition data sets.

If you delete and redefine VSAM data sets, you receive a z/OS IEC161I system message when reloading a partition. This is not an error message. It indicates that a VSAM data set was empty when it was opened. Figure 215 shows the message for an ILDS.

```
IEC161I 152-061,JOUK03D,RELOAD,PE001L,,
IEC161I JOUK03.HALDB.DB.PEOPLE.L00001,
IEC161I JOUK03.HALDB.DB.PEOPLE.L00001.DATA,CATALOG.TOTICF2.VTOTCAT
```

Figure 215. Example: IEC161I message during reload

Related Reading: For more information on IEC system messages, see *z/OS V1R4: MVS System Messages, Vol 7 (IEB-IEE)*.

Reloading HALDB Partitions and Databases for Offline Reorganization

HALDB partitions and databases can be reloaded with the HD Reorganization Reload utility (DFSURGL0). The HD Reorganization Reload utility reads the output file from the HD Reorganization Unload utility. You do not specify the partitions to be reloaded. They are determined by the records in the input file to the HD Reorganization Reload utility.

Do not include DD statements for the HALDB database data sets. The HD Reorganization Reload utility uses dynamic allocation for HALDB data sets. This is not true for non-HALDB databases.

You must supply buffer pools for all data sets in the partitions that are reloaded. This includes the ILDSs.

The HD Reorganization Reload utility sets the image copy needed flag for data sets in partitions that it loads. You should image copy them as you would any database data sets after they have been reloaded.

Figure 216 shows a sample job that reloads HALDB partitions. The partitions it reloads depend on the records in the input file.

```
//JOUK03D JOB (999,POK),JOUK03,CLASS=A,NOTIFY=&SYSUID,
//          MSGLEVEL=(1,1),MSGCLASS=X,REGION=0M
//JOB LIB DD DSN=IMSPSA.IMS0.SDFSRESL,DISP=SHR
//          DD DSN=IMSPSA.IM0A.MDALIB,DISP=SHR
//*****
//* HD RELOAD FOR THE PEOPLE DATABASE
//*****
//RELOAD EXEC PGM=DFSRR00,REGION=1024K,
//          PARM='ULU,DFSURGL0,PEOPLE,,,,,,Y,N'
//DFSRESLB DD DSN=IMSPSA.IMS0.SDFSRESL,DISP=SHR
//IMS DD DISP=SHR,DSN=JOUK03.HALDB.DBDLIB
//DFSUINPT DD DSN=JOUK03.UNLOAD.PEOPLE,DISP=OLD
//DFSVSAMP DD *
VSRBF=8192,50
IOBF=(4096,50)
/*
//SYSPRINT DD SYSOUT=*
//DFSSTAT DD SYSOUT=*
```

Figure 216. Example: JCL to Reload a HALDB Partition

ILDS Reorganization Updates: The HD Reorganization Reload utility updates the ILDS for partitions that contain targets of logical relationships or secondary indexes. The utility has three options for updating ILDSs:

- No control statement
- NOILDS control statement
- ILDSMULTI control statement

If you do not specify a control statement in the SYSIN data for the HD Reorganization Reload utility, an ILDS entry is updated or created when a target of a secondary index or logical relationship is inserted in the partition. An entry exists if a previous reorganization loaded the target segment in the partition. The updates to the ILDS are done in VSAM update mode. When a CI or CA is filled, it must be split by VSAM. Free space in the ILDS can help avoid these splits. Updates can be random or sequential. This depends on the order in which these segments are inserted and their ILKs. The ILDS keys are based on the ILK that is based on the location of the target segment when it was created.

Reorganizing HALDBs

You can create free space in an ILDS by copying it using the VSAM REPRO command. The REPRO command honors the free space parameters in the VSAM DEFINE.

You can delete and redefine the ILDS before reloading. You might want to do this to eliminate entries in the ILDS for target segments that are no longer in the partition. The HD Reorganization Reload utility never deletes an entry in the ILDS. The only way to delete these entries is to delete and redefine the ILDS. Alternately, an empty ILDS contains no free space. A reload with a large number of target segments might require a large number of CI and CA splits.

If you specify a NOILDS control statement in the SYSIN data, the HD Reorganization Reload utility does not update or create entries in the ILDSs. They must be created by a separate process using the HALDB Index/ILDS Rebuild utility (DFSPREC0). Separate executions of DFSPREC0 are required for each partition. These executions can be done in parallel and on different machines.

The ILDSMULTI option applies only to migration reloads. For more information about ILDSMULTI, see the HD Reorganization Reload utility section of *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

Reorganizing HALDB Partitioned Secondary Index Databases: You might need to reorganize your partitioned secondary index (PSINDEX) database. Because the reorganization of HALDBs does not require the recreation of their secondary indexes, a PSINDEX database can become disorganized as entries are added to it over time.

The HD Reorganization Unload utility and the HD Reorganization Reload utility can be used to reorganize PSINDEX databases. The restrictions and recommendations for reorganizing other HALDB databases also apply to PSINDEX databases with one exception: HALDB secondary indexes have no ILDSs. The HD Reorganization Reload utility control statements should not be used with secondary indexes.

The steps for reorganizing a PSINDEX database are the same as those for reorganizing other types of HALDBs offline. See “Overview of HALDB Offline Reorganization” on page 365 for a list of these steps.

HALDB Online Reorganization

Prior to IMS Version 9, you had to ensure that HALDB partitions were offline before you could perform database reorganization for them. IMS Version 9 introduced an integrated HALDB Online Reorganization function that allows HALDB partitions to remain online and available for IMS application programs during a database reorganization.

An online reorganization of a PHDAM or PHIDAM HALDB partition runs non-disruptively, allowing concurrent IMS updates, including updates by data-sharing IMS systems. The online reorganization is non-disruptive because IMS copies small amounts of data from the partition's original data sets (the input data sets) to separate output data sets. IMS tracks which data has been copied so that IMS applications can automatically retrieve or update data from the correct set of data sets.

As described in “Data Set Naming Conventions for HALDB Online Reorganization” on page 378, HALDB Online Reorganization extends the established data definition and data set naming convention for HALDBs. The data set groups in a HALDB database use the characters A-through-J in the DDNAMEs and data set names of

the ten supported data set groups. The primary index for a PHIDAM database uses the character X in these names. This data definition and data set naming convention is extended so that IMS uses the characters M-through-V (and Y) for an alternate set of data sets.

The initial load or offline reorganization reload of a HALDB partition always uses the A-through-J (and X) data sets. Until the first time that you reorganize a HALDB partition online, only the A-through-J (and X) data sets are used.

There are three phases of online reorganization for a HALDB partition:

1. The initialization phase, during which IMS prepares the output data sets and updates the RECON data sets.
2. The copying phase, during which IMS performs the actual reorganization by copying the data from the input data sets to the output data sets.
3. The termination phase, during which IMS closes the input data sets and updates the RECON data sets.

The Initialization Phase for HALDB Online Reorganization

You start the online reorganization of a HALDB partition using the `INITIATE OLREORG` command. See the *IMS Version 9: Command Reference* for more information about this command.

During the initialization phase, IMS updates the RECON data sets to establish the ownership of the online reorganization by the IMS system that is performing the online reorganization. This ownership means that no other IMS system can perform a reorganization of the HALDB partition until the current online reorganization is complete or until ownership is transferred to another IMS system. IMS adds the M-V (and Y) DBDSs to the RECON data sets if those DBDS records do not already exist. IMS also adds the M-V (and Y) DBDSs to any existing change accumulation groups and DBDS groups that include the corresponding A-J (and X) DBDSs.

Before online reorganization begins for a HALDB partition, there is a single set of *active* data sets for the HALDB partition. These active data sets are the input data sets for the copying phase. There might also be a set of *inactive* data sets from a prior online reorganization that are not used by IMS application programs.

During the initialization phase, IMS evaluates each of the inactive data sets to ensure that it meets the requirements for output data sets (see “HALDB Online Reorganization Requirements for Existing Output Data Sets” on page 551). If any of the output data sets does not exist, IMS creates it automatically during this phase.

At the end of the initialization phase, IMS treats the original active set of data sets as the input set and the inactive data sets as the output set. This use of the input and output sets of data sets is represented by the *cursor-active* status for the partition, which is recorded in online reorganization records in the RECON data sets. A listing of the partition's database record in the RECON data sets shows `OLREORG CURSOR ACTIVE=YES`. A listing of the partition also shows that both sets of DBDSs are active: the first set of DBDSs listed is for the input data set and the second set of DBDSs is for the output data set, for example, `DBDS ACTIVE=A-J` and `M-V`. While the partition is in the cursor-active status, both sets of data sets must be available for the partition to be processed by any application.

Figure 217 on page 372 shows part of a listing of the RECON data sets for a HALDB partition that has the cursor-active status.

Figure 217. Example RECON Listing: DB Record for a HALDB in Cursor-Active Status

Also during the initialization phase of an online reorganization, IMS dynamically creates a program specification block (PSB) for the reorganization. The PSB name is the 7-character HALDB partition name prefixed with the single character zero ('0'). For example, a HALDB partition with the name SSN5603 has a dynamic PSB with the name 0SSN5603 for the reorganization work. This PSB does not exist in the PSBLIB or the ACBLIB, but the name can appear in a listing of RECON or in output from utilities.

During the copying phase, the HALDB partition comprises the A-through-J (and X) data sets and the M-through-V (and Y) data sets. During this phase, both sets of data sets must be available in order for IMS applications to access the partition.

372 Administration Guide: Database Manager

the application program. Figure 218 illustrates the relationship between the input and output data sets at a point during the online reorganization.

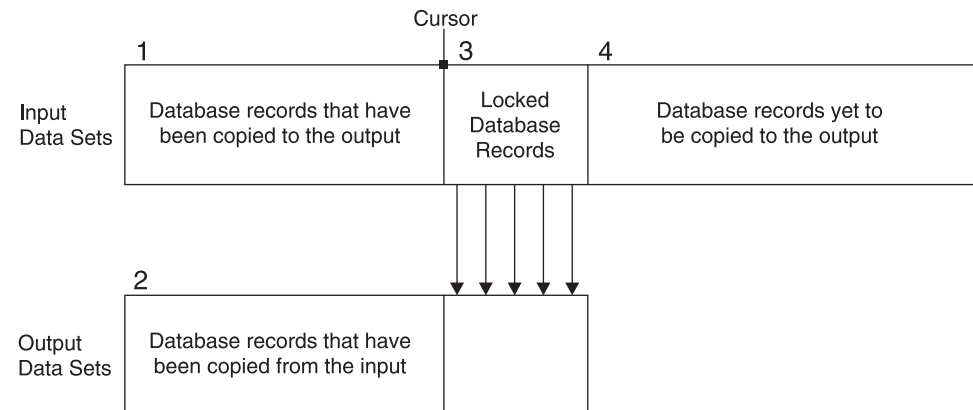


Figure 218. The Relationship between Input Data Sets and Output Data Sets during the Online Reorganization of a HALDB Partition

Figure 218 shows two sets of database data sets for a HALDB partition, the input data sets that have not been reorganized and the output data sets that have been (at least partially) reorganized. The figure shows the reorganization as progressing from left to right, from the input data sets above to the output data sets below. The data sets in the figure are divided into four areas:

1. Data within the input data sets that has been copied to the output data sets. This area reflects the old data organization (prior to the reorganization), and is not used again by IMS applications until the data sets are reused as the output data sets for a later online reorganization.
2. Data within the output data sets that has been copied from the input data sets. This data in this area has been reorganized, and can be used by IMS applications during the reorganization.
3. Data within both the input and output data sets that is locked and in the process of being copied and reorganized from the input data sets to the output data sets. This area of locked records is called a *unit of reorganization*. From a recovery point of view, this unit of reorganization is equivalent to a unit of recovery.

While IMS processes the current unit of reorganization, IMS applications that access any of the locked data records must wait until IMS completes the reorganization for those records. After the copying and reorganization completes for the unit of reorganization, IMS commits the changes and unlocks the records, thus making them available again for IMS applications.
4. Data within the input data sets that has not yet been copied to the output data sets. This data has also not yet been reorganized, and can be used by IMS applications during the reorganization.

As the online reorganization progresses, IMS uses a kind of pointer called a *cursor* to mark the end point of those database records that have already been copied from the input data sets to the output data sets. As the reorganization and copying proceeds, this cursor moves through the partition (from left to right in Figure 218).

When an IMS application program accesses data from a HALDB partition that is being reorganized online, IMS retrieves the data record:

Reorganizing HALDBs

- From the output data sets if the database record is located “at or before” the cursor.
- From the input data sets if the database record is located “after” the cursor.

If the data record happens to fall within the unit of reorganization, IMS retries the data access after the records are unlocked. An application program does not receive an error status code for data within a unit of reorganization.

To allow recovery of either an input data set or an output data set, all database changes are logged during the online reorganization, including the database records that are copied from the input data set to the output data sets.

The Termination Phase for HALDB Online Reorganization

The online reorganization of a HALDB partition terminates after the end of the copying phase, or when IMS encounters an error condition during the reorganization. You can also stop the online reorganization of a HALDB partition using the `TERMINATE OLREORG` command. See the *IMS Version 9: Command Reference* for more information about this command.

After the copying phase is complete for a HALDB partition, the output data sets become the *active* data sets, and the input data sets become the *inactive* data sets. The active data sets are used for all data access by IMS application programs. The inactive data sets are not used by application programs, but can be reused for a subsequent online reorganization. Unless you perform an initial load or a batch reorganization reload for the partition, successive online reorganizations for the partition alternate between these two sets of data sets.

IMS updates the partition’s database record in the RECON data sets to reset the cursor-active status for the partition to reflect that there is now just one set of data sets. A listing of this record from the RECON data sets shows `OLREORG CURSOR ACTIVE=NO` and the `ACTIVE DBDS` field shows the active (newly reorganized) data sets. IMS also updates the online reorganization records in the RECON data sets with the timestamp of when the reorganization completed.

If you specified the `DEL` keyword for the `INITIATE OLREORG` command (or the `UPDATE OLREORG` command), IMS deletes the inactive data sets after resetting the cursor-active status for the partition. Before deleting the inactive data sets, IMS notifies all sharing IMS systems, including batch jobs, that the online reorganization is complete and is recorded in the RECON data sets. The IMS system that is performing the online reorganization waits until it receives an acknowledgement from each of these sharing IMS systems that they have closed and deallocated the now-inactive data sets, and then it deletes these data sets. However, if the acknowledgements are not received within 4.5 minutes, the owning IMS system will attempt to delete the inactive data sets anyway. Without the acknowledgements, the deletion attempt is likely to fail.

Finally, at the end of the termination phase, IMS updates the RECON data sets to reset the ownership of the online reorganization so that no IMS system has ownership. This resetting of ownership means that any IMS system can perform a subsequent reorganization of the HALDB.

If online reorganization of a HALDB partition terminates prior to completion, either because of an error or because you issued the `TERMINATE OLREORG` command, you must restart the online reorganization or perform an offline reorganization for the partition.

Figure 219 shows the normal processing steps of a successful online reorganization of a HALDB partition. The columns represent the flow of control through the phases of the online reorganization, from the user to IMS, and the status of the data sets as the processing events occur.

| Phase | User | IMS | Status of Data Sets |
|----------------|-------------------------------------|---|--|
| Initialization | Invokes INITIATE OLREORG command | Gets ownership from DBRC | |
| | | Performs Init. steps (create data sets, etc.) | |
| | | Tells DBRC about cursor-active status | Now two active sets of data sets exists OLREORG CURSOR ACTIVE=YES in RECON listing |
| Copying | | Copies data from input to output data sets | Database records are locked, copied and then unlocked by UOR |
| Termination | | When copying is complete, tells DBRC to end cursor-active status | Now only one active set of data sets exists OLREORG CURSOR ACTIVE=NO in RECON listing |
| | | Possibly deletes inactive set of data sets | |
| | | Relinquishes ownership through DBRC | |

Figure 219. Normal Processing Steps of HALDB Online Reorganization

Migration and Coexistence Considerations for HALDB Online Reorganization

During migration from a prior IMS release to IMS Version 9, IMS marks the DB records in the RECON data sets for all existing HALDBs to indicate that they cannot be reorganized online. To allow online reorganization of a HALDB partition, use the DBRC `CHANGE.DB DBD(HALDB_master) OLRCAP` command to change the status of the HALDB and all of its partitions to allow online reorganization for those partitions. You can also use the `CHANGE.DB ALL OLRCAP` command to enable online reorganization for all of your HALDB databases.

Table 27 on page 376 shows the IMS versions that can access HALDBs that are capable of being reorganized online.

Reorganizing HALDBs

Table 27. IMS Versions that Can Access HALDBs that Are Capable of Being Reorganized Online

| IMS Version | Access to HALDB partitions that are capable of being reorganized online? |
|--|--|
| IMS Version 7 | No |
| IMS Version 8 | No |
| IMS Version 8 with the OLR Coexistence SPE | Yes |
| IMS Version 9 | Yes |

You must apply the IMS Version 8 OLR Coexistence SPE to allow full data sharing between IMS Version 8 and IMS Version 9 systems that have HALDBs that are capable of being reorganized online.

You must use the following IMS Version 9 (or later) utilities to process HALDBs that are capable of being reorganized online:

- Database Recovery
- Database Image Copy
- Database Image Copy 2
- Database Change Accumulation

Fallback Considerations for HALDB Online Reorganization

For IMS Version 8 systems that have the OLR Coexistence SPE applied, you can access or share data with all Version 9 HALDBs. However, if any M-through-V (and Y) data sets are active, or if the HALDB Online Reorganization cursor is active for any partitions, you must use IMS Version 9 utilities whenever those partitions are processed, except for the Batch Backout, Log Recovery, and Log Archive utilities which must be run on the release of IMS that created the logs.

For any partitions with M-through-V (and Y) data sets active, or for any partitions with an active HALDB Online Reorganization cursor, you must run an offline reorganization before you can fall back to using the IMS Version 8 utilities.

Should fallback to a prior version become necessary, you must define all the HALDBs as no longer capable of being reorganized online. For IMS Version 7 systems and IMS Version 8 systems that do not have the OLR Coexistence SPE applied, you can access only those HALDBs that are not capable of being reorganized online. After fallback, HALDBs that are capable of being reorganized online are unavailable until you complete the following actions:

1. Using the IMS Version 9 offline reorganization utility, reorganize all partitions that have the M-through-V (and Y) data sets active; these data sets could be active either because the partition has the cursor-active status or because these are the only data sets for the partition.
2. Define the partitions as no longer capable of being reorganized online by using the command `CHANGE.DB DBD(HALDB_master) OLRNOCAP`.

Restrictions for HALDB Online Reorganization

The following restrictions apply to HALDB Online Reorganization:

- You can perform an online reorganization only for a HALDB that is defined in the RECON data sets as capable of being reorganized online (OLRCAP). For more

information about the OLRCAP parameter, see the INIT.DB command or the CHANGE.DB command in the *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

- You cannot start an online reorganization for a partition if another IMS system already owns an online reorganization for that partition.
- You cannot make data definitional changes during an online reorganization of a partition. HALDB Online Reorganization provides only reclustering and space distribution advantages.
- Image copy for a partition is not allowed if the partition is in the cursor-active status. This restriction applies even if the online reorganization terminated before the cursor-active status has been reset and the online reorganization for the partition is not owned by any IMS.
- To backout in-flight work from an online reorganization, you must run a batch backout using a DL/I region type.
- To use a type-2 command to start an online reorganization for a HALDB partition, you must have an IMS Common Service Layer that includes the Operations Manager and the Structured Call Interface. See the *IMS Version 9: Common Service Layer Guide and Reference* for more information.
- HALDB Online Reorganization runs only in a local storage option-subordinate (LSO=S) environment. IMS rejects attempts to initiate an online reorganization for a HALDB partition in a local storage option-yes (LSO=Y) environments. For more information about the LSO specification, see the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.
- You cannot perform an online reorganization for a HALDB partition from an alternate IMS system in an XRF complex. However, after an XRF takeover, the new active IMS system will continue a reorganization that was active when the takeover process began.
- You cannot perform an online reorganization for a HALDB partition from a tracking IMS system in an RSR complex. However, for HALDBs that are registered as DBTRACK at the tracking IMS system, IMS tracks the effects of an online reorganization in the same way it tracks updates to any database. See “IMS Remote Site Recovery Processing for HALDB Online Reorganization” on page 384 for more information.
- You cannot issue the following commands for a HALDB partition while it is being reorganized online:

- /START DATABASE or UPDATE DATABASE NAME(*name*) START(ACCESS)
- /DBRECOVERY DATABASE or UPDATE DATABASE NAME(*name*) STOP(ACCESS)
- /DBDUMP DATABASE or UPDATE DATABASE NAME(*name*) STOP(UPDATES)
- /STOP DATABASE or UPDATE DATABASE NAME(*name*) STOP(SCHD)

If you issue any of these commands for a HALDB partition that is actively being reorganized online, IMS displays error message DFS0488I and does not process the command for the named partition. For more information about these commands, see the *IMS Version 9: Command Reference*. For more information about message DFS0488I, see the *IMS Version 9: Messages and Codes, Volume 2*.

- You cannot issue the following commands for a HALDB master while any of its partitions is being reorganized online:
- /START DATABASE ACCESS UP or UPDATE DATABASE NAME(*name*) START(ACCESS)
- /DBRECOVERY DATABASE or UPDATE DATABASE NAME(*name*) STOP(ACCESS)
- /DBDUMP DATABASE or UPDATE DATABASE NAME(*name*) STOP(UPDATES)

For more information about these commands, see the *IMS Version 9: Command Reference*.

Data Set Naming Conventions for HALDB Online Reorganization

As described in “HALDB Naming Conventions” on page 22, the data sets for HALDB partitions use a specified naming convention. HALDB Online Reorganization extends this naming convention to include a second set of data sets. Data sets for partitions that are involved in HALDB Online Reorganization use the following naming convention: *bbbbbbb.dppppp*

| | |
|----------------|---|
| bbbbbbb | Represents a data set name prefix of up to 37 characters that you defined using the HALDB Partition Definition utility or DBRC batch command (INIT.DB, INIT.PART, CHANGE.DB, or CHANGE.PART). The same data set base name is used for every data set within a HALDB partition. |
| d | Represents an IMS-assigned data set name type character that uniquely identifies a specific data set for a HALDB partition. The possible single-character values are: A-through-J “A” corresponds to the first, or possibly only, data set group that is defined in the DBD, “B” corresponds to the second data set group, and so on. The use of the characters A-through-J applies generally to HALDB partitions, regardless of whether they are capable of being reorganized online. M-through-V “M” corresponds to the first, or possibly only, data set group that is defined in the DBD, “N” corresponds to the second data set group, and so on. The use of the characters M-through-V applies only to HALDB partitions that are capable of being reorganized online. L The indirect list data set (ILDS). The online reorganization process does not make a copy of this data set. X The primary index of a PHIDAM database. This data set is the index for the A-through-J data sets and is replaced by the Y data set when the online reorganization process copies the database records from the A-through-J and X data sets into the M-through-V and Y data sets. The use of the X character applies generally to HALDB partitions, regardless of whether they are capable of being reorganized online. Y The primary index of a PHIDAM database. This data set is the index for the M-through-V data sets and it is replaced by the X data set when the online reorganization process copies the database records from the M-through-V and Y data sets into the A-through-J and X data sets. The use of the Y character applies only to HALDB partitions that are capable of being reorganized online. |
| ppppp | Specifies the five-digit partition ID that is assigned by IMS. |

The data set names for the output data sets are identical to the names of the corresponding input data sets, except for the IMS-assigned data set name type character (A-through-J, M-through-V, X, or Y). Table 28 shows example data set names.

Table 28. Data Set Name Examples for HALDB Online Reorganization

| Active Data Set Before Online Reorganization | Data Set Group or Index | Partition ID | Input Data Set Name | Output Data Set Name |
|--|-------------------------|--------------|---------------------|----------------------|
| A-through-J (and X) | 1 | 00003 | DH41.A00003 | DH41.M00003 |
| A-through-J (and X) | Index | 00065 | ACCT.X00065 | ACCT.Y00065 |
| M-through-V (and Y) | 2 | 00005 | PAY.MST.N00005 | PAY.MST.B00005 |
| M-through-V (and Y) | 8 | 00001 | PAY.EMP.T00001 | PAY.EMP.H00001 |

Output Data Set Requirements for HALDB Online Reorganization

During the initialization phase for an online reorganization of a HALDB partition, IMS creates any output data sets that do not already exist. For example, when the input data sets are the A-through-J (and X) set, if the M and P output data sets already exist, but the N and O output data sets do not, IMS creates the N and O data sets and uses the existing M and P data sets.

Any existing output data sets must have the characteristics described in “HALDB Online Reorganization Requirements for Existing Output Data Sets” on page 551. Any data in the existing output data sets is overwritten during the copying phase of an online reorganization. Output data sets that IMS creates for the online reorganization have the characteristics described in “Attributes of Automatically-Created Output Data Sets” on page 551.

Starting HALDB Online Reorganization

Table 29 describes the tasks and commands for starting or resuming an online reorganization for a HALDB partition.

Table 29. Mapping Startup Tasks to Commands for HALDB Online Reorganization

| Task | Command | Command Type |
|---|--|--------------|
| Specify that a HALDB master is capable of being reorganized online. | INIT.DB OLRCAP or CHANGE.DB OLRCAP | DBRC |
| Begin HALDB Online Reorganization for one or more partitions. | INITIATE OLREORG | Type 2 |
| Begin HALDB Online Reorganization for one or more partitions. | /INITIATE OLREORG | Type 1 |
| Resume HALDB Online Reorganization for one or more partitions. | INITIATE OLREORG | Type 2 |
| Resume HALDB Online Reorganization for one or more partitions. | /INITIATE OLREORG | Type 1 |
| Set the RATE for a HALDB Online Reorganization. | INITIATE OLREORG SET(RATE(rate)) or UPDATE OLREORG SET(RATE(rate)) | Type 2 |
| Set the RATE for a HALDB Online Reorganization. | /INITIATE OLREORG SET(RATE(rate)) or /UPDATE OLREORG SET(RATE(rate)) | Type 1 |

Related Reading:

Reorganizing HALDBs

- For more information about the INITIATE OLREORG, /INITIATE OLREORG, UPDATE OLREORG and /UPDATE OLREORG commands, see the *IMS Version 9: Command Reference*.
- For more information about the CHANGE.DB and the INIT.DB commands, see the *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

Monitoring HALDB Online Reorganization

Table 30 describes the tasks and commands for monitoring an online reorganization for a HALDB partition.

Table 30. Mapping Monitoring Tasks to Commands for HALDB Online Reorganization

| Task | Command | Command Type |
|--|----------------------|--------------|
| Display status and rate information about HALDB Online Reorganizations that are in progress. | QUERY OLREORG | Type 2 |
| Monitor and display the status of the specified databases or partitions (including those HALDB Online Reorganizations that are in progress). | /DISPLAY DB OLR | Type 1 |
| Display HALDB Online Reorganization status. | QUERY DB | Type 2 |
| List all of the databases for which HALDB Online Reorganizations are in progress. | QUERY DB STATUS(OLR) | Type 2 |

Related Reading: For more information about the /DISPLAY DB, /DISPLAY DB OLR, QUERY DB, and QUERY OLREORG commands, see the *IMS Version 9: Command Reference*.

Modifying and Tuning HALDB Online Reorganization

Table 31 describes the tasks and commands for modifying and tuning an online reorganization for a HALDB partition.

Table 31. Mapping Modifying and Tuning Tasks to Commands for HALDB Online Reorganization

| Task | Command | Command Type |
|---|---|--------------|
| Stop HALDB Online Reorganization for one or more partitions. | TERMINATE OLREORG | Type 2 |
| Stop HALDB Online Reorganization for one or more partitions. | /TERMINATE OLREORG | Type 1 |
| Change the impact of HALDB Online Reorganization on overall system performance, for one or more partitions. | UPDATE OLREORG SET(RATE(<i>rate</i>)) | Type 2 |
| Change the impact of HALDB Online Reorganization on overall system performance, for one or more partitions. | /UPDATE OLREORG SET(RATE(<i>rate</i>)) | Type 1 |
| Specify whether to delete the inactive data sets after the copying phase completes. | UPDATE OLREORG OPTION(DEL NODEL) | Type 2 |
| Specify whether to delete the inactive data sets after the copying phase completes. | /UPDATE OLREORG OPTION(DEL NODEL) | Type 1 |

Related Reading: For more information about the TERMINATE OLREORG, /TERMINATE OLREORG, UPDATE OLREORG, and /UPDATE OLREORG commands, see the *IMS Version 9: Command Reference*.

Example: Figure 220 on page 382 shows the processing steps for an online reorganization of a HALDB partition and how it is affected by a TERMINATE OLREORG command that temporarily stops the reorganization:

- When you issue the TERMINATE OLREORG command, IMS terminates the reorganization by entering the termination phase.
- Later, when you issue the INITIATE OLREORG command, IMS restarts the reorganization from the initialization phase, then proceeds to the copying phase. In the figure, the reorganization then completes successfully through the termination phase.

Note that there are two sets of data sets for the second initialization phase because the reorganization is not complete.

In the figure, the columns represent the flow of control through the phases of the online reorganization, from the user to IMS, and the status of the data sets as the processing events occur.

Reorganizing HALDBs

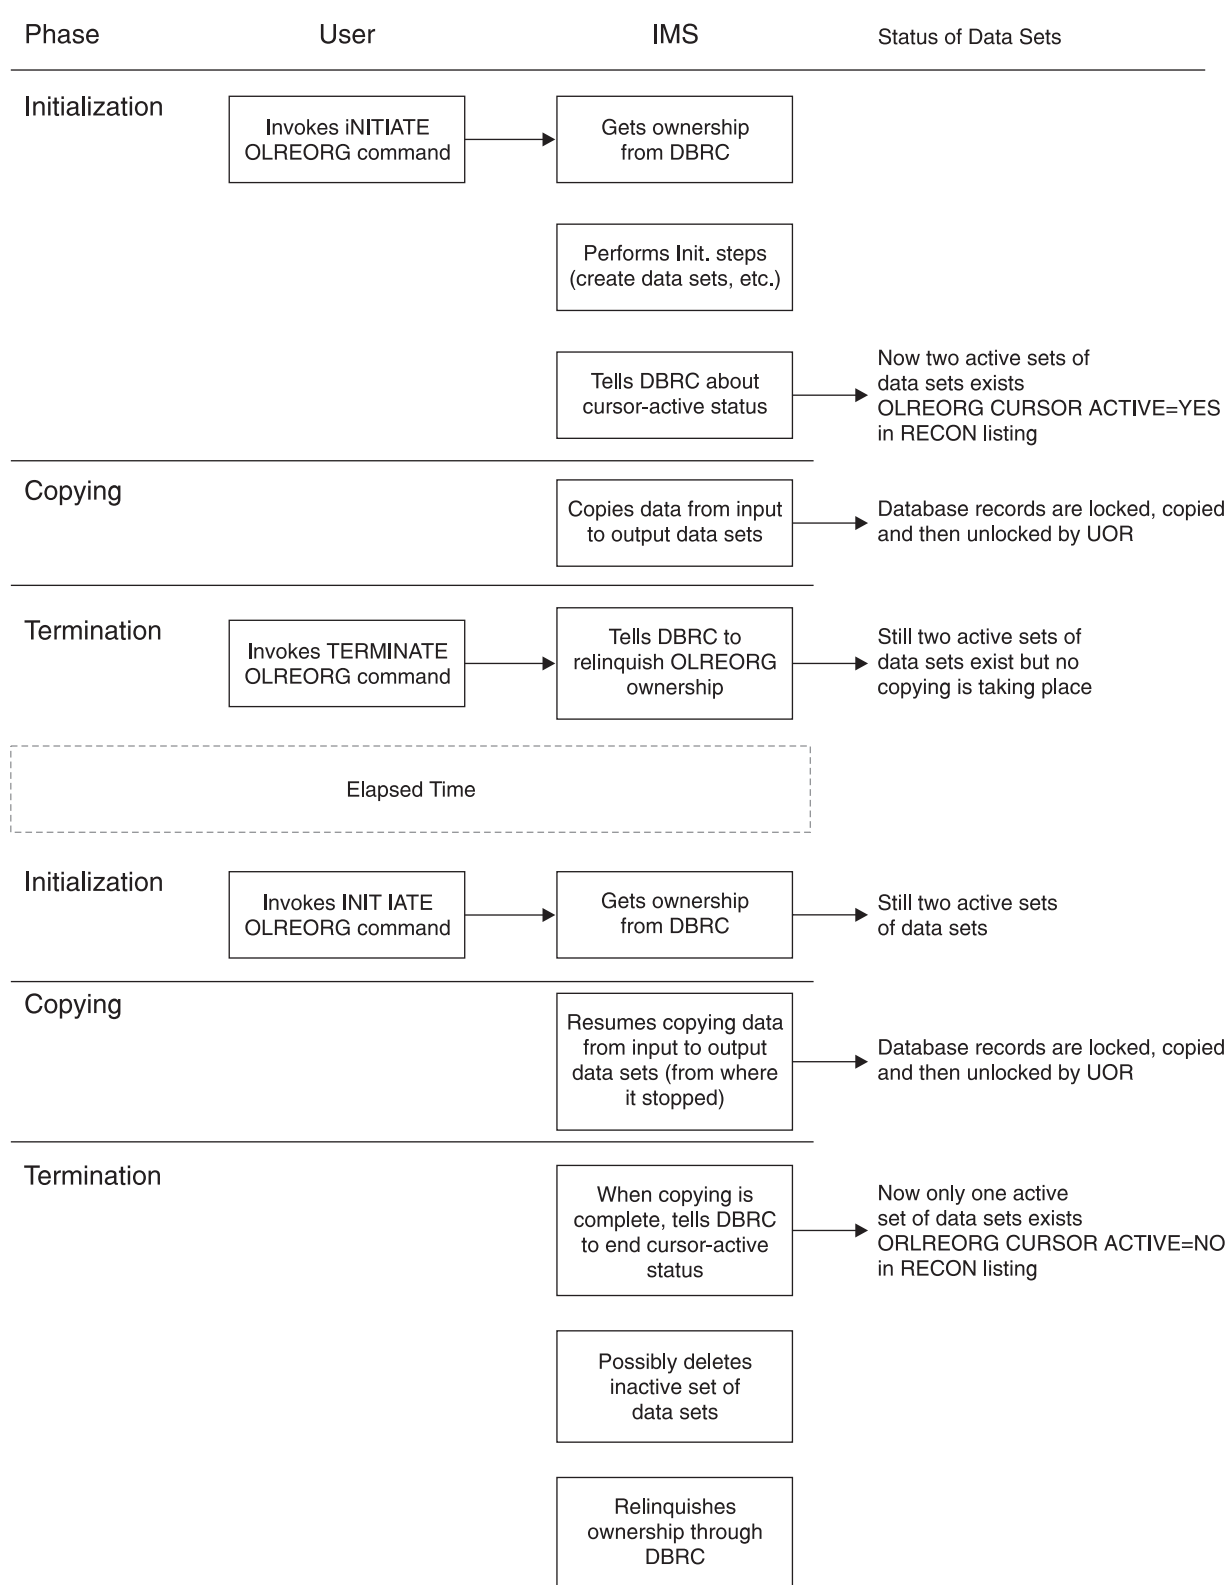


Figure 220. Processing Steps for an Interrupted Online Reorganization of a HALDB Partition

How a HALDB Online Reorganization Impacts IMS Logging

The online reorganization of a HALDB partition generates X'50' database change log records for all of the data in the partition. IMS also logs other HALDB Online Reorganization information in a small number of X'29' log records. Thus, the total amount of log data generated by a HALDB Online Reorganization is considerably greater than the amount of data in the partition.

Reorganizing partitions online, especially reorganizing multiple partitions in parallel on the same IMS system, can generate sufficient log data to impact normal transaction processing. The large number of log records that are generated can also affect the rate of OLDS switches and log archiving.

Controlling the Overall System Impact of a HALDB Online Reorganization

An online reorganization of a HALDB partition can impact the overall system performance of an IMS system, and likewise, other IMS work can affect the performance of an online reorganization. You can use the RATE parameter of the INITIATE OLREORG and UPDATE OLREORG commands to control the impact of an online reorganization on your IMS system.

Depending on system resources, IMS requires a certain amount of time to process a unit of reorganization. The value of the RATE parameter represents how much of that time IMS spends actually copying and reorganizing the data and how much of that time IMS spends in an intentionally introduced delay. You specify the RATE value as a percentage, with values less than 100 representing the addition of an intentionally introduced delay. Adding this delay to the copying process can help minimize the online reorganization's impact on other IMS work.

The default value for the RATE parameter is 100, which allows the online reorganization to run as fast as possible, depending on system resources, system contention, and log contention, with no intentionally introduced delay. However, if you set the RATE value to 25, for example, IMS adds a delay to the reorganization processing so that 25% of the total processing time for a unit of reorganization is spent copying the data, and the remaining 75% is spent in an intentionally introduced delay. Thus, RATE(25) would cause the online reorganization to take approximately four times as long to run as it would have run with RATE(100).

You can change the RATE value at any time by issuing the UPDATE OLREORG command.

IMS Restart and XRF Processing for HALDB Online Reorganization

If you shut down IMS while any online reorganizations of HALDB partitions are running, IMS suspends the reorganizations before completing the shutdown checkpoint. After IMS restarts, IMS automatically resumes the online reorganizations. IMS resumes the online reorganization even if you specified the NOPDBO option in the DFSVSMxx member.

If IMS terminates abnormally while any online reorganizations are running, IMS dynamically backs out all uncommitted changes for these reorganizations to the most recent sync point. After IMS restarts, IMS automatically resumes the online reorganizations.

Likewise, when an XRF takeover occurs, IMS automatically resumes the online reorganizations on the new active IMS system.

IMS Restart and Fast Database Recovery Processing for HALDB Online Reorganization

When an FDBR takeover occurs, IMS performs the following actions:

- Backs out all uncommitted changes for these reorganizations to the most recent sync point
- Closes the partition
- Unauthorizes the partition

When you restart the IMS system, IMS does not resume the online reorganization because the partitions are not authorized after the FDBR terminates.

IMS Remote Site Recovery Processing for HALDB Online Reorganization

During Remote Site Recovery (RSR) tracking, IMS updates the tracking site RECON data sets with information for HALDB Online Reorganization. For HALDB partitions that are registered as database-level tracking (DBTRACK) at a Database Level Tracking (DLT) tracking IMS system, IMS performs the following steps during tracking of the online reorganization:

- Creates the output data sets for the shadow partition, as needed.
- Updates both the input and output data sets for the shadow partition.
- Marks the original input data sets as inactive, and marks the output data sets as the active data sets at the completion of the tracking of the online reorganization.
- Deletes the inactive data sets if delete option is in effect. You specify this option (or accept the default) by using the OPTION keyword of the INITIATE OLREORG command at the active site.

IMS stops the shadow partition if errors occur during the validation or creation of the output data sets. The tracked partition at the active site is unaffected by errors at the tracking site. After you correct the problem that caused the error, restart the shadow partition on the tracking IMS system to initiate online forward recovery for the partition and to continue tracking.

If the output data sets for the online reorganization already exist at the tracking site before tracking begins, ensure that these data sets have same characteristics (such as block size, record size, and control interval size) as those at the active site. See “HALDB Online Reorganization Requirements for Existing Output Data Sets” on page 551 for the data set characteristics. If you change output data set characteristics manually at the active site, you must make the same changes at the tracking site.

After an RSR takeover, IMS stops all HALDB partitions, including those that had online reorganizations in process. After you rebuild the primary index and indirect list data sets using the HALDB Index/ILDS Rebuild utility (DFSPREC0) at the new active site, issue the INITIATE OLREORG command to resume the online reorganizations, if needed. The online reorganizations are not automatically restarted after takeover.

Locking Impacts of HALDB Online Reorganization

An online reorganization for a HALDB partition uses the IRLM to request global locks for each UOR to maintain partition integrity and recoverability. An online reorganization that is running with the RATE(100) specification can incur a significant number of IRLM lock structure accesses per second.

Recommendations:

- Consider using a second subpool to relieve database buffer contention for more than four concurrent online reorganizations.
- Use the IBM CFSizer to model the additional coupling facility activities to ensure that your coupling facility configuration is capable of handling the extra load introduced by the online reorganizations:
 - For IRLM 2.1 with PC=N0 specified, each additional 1000 concurrently held locks requires 256 KB of ECSA storage.
 - For IRLM 2.2, each additional 1000 concurrently held locks requires 540 KB obtained from IRLM private storage. No increase in ECSA storage is necessary.
- Review your LOGL latch contention rate, OLDS logging rate, IRLM lock structure access, and DBBP (for OSAM) latch contention.

Using IMS Utilities with HALDB Online Reorganization

The following IMS utilities have special considerations when used with HALDB Online Reorganization: Batch Backout (DFSBB000), Database Change Accumulation (DFSUCUM0), Database Image Copy (DFSUDMP0), Database Recovery (DFSURDB0), and Primary Index and ILDS Rebuild (DFSPRECO).

Batch Backout (DFSBB000)

If dynamic backout fails for a unit or reorganization, IMS creates a backout record in the RECON data sets that contains the dynamic PSB name. Run the Batch Backout utility (DFSBB000) for the listed PSB name.

If dynamic backout was not attempted, use the Log Recovery utility (DFSULTR0) with the PSB option to list those PSBs that require backout. If there is in-flight online reorganization work for a HALDB partition that requires backout, the Log Recovery utility lists the dynamic PSB name. Run the Batch Backout utility (DFSBB000) for each of the listed PSB names.

Database Change Accumulation (DFSUCUM0)

You can use the Database Change Accumulation utility (DFSUCUM0) to accumulate changes for HALDB partition A-through-J data sets and for the M-through-V data sets. You need to specify the DB0 control statement so that the Database Change Accumulation utility accumulates changes or purges changes from before the online reorganization started.

The Database Change Accumulation utility might create Database Change Accumulation header records (type X'25' records) for a corresponding A-through-J and M-through-V data set if the online reorganization checkpoint is not complete at the start of the database change accumulation.

Database Image Copy (DFSUDMP0)

You can use the Database Image Copy utility (DFSUDMP0) to copy the currently active data set that is recorded in the RECON data sets. The Database Image Copy utility also determines if it should copy the M-through-V data sets or the A-through-J data sets. However, if a partition is in the cursor-active status, you cannot run the Database Image Copy utility for that partition.

It is not necessary to code a DD statement in the JCL when copying HALDB partition data sets, because they are dynamically allocated.

Database Recovery (DFSURDB0)

The Database Recovery utility (DFSURDB0) expects the utility output data sets to exist, and makes no attempt to create them. See “Recovery for HALDB Online Reorganization” on page 386.

Primary Index and ILDS Rebuild (DFSPREC0)

You can use the HALDB Index/ILDS Rebuild utility (DFSPREC0) to recover the primary index data set for both the input and output data sets. You can also use the HALDB Index/ILDS Rebuild utility to recover both the X and Y primary index data sets and the ILDS in a single run. You cannot specify a particular input or output data set, but if the partition is in the cursor-active status, the utility allocates and rebuilds all of the necessary data sets.

Related Reading: For more information about these utilities, see the *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

Recovery for HALDB Online Reorganization

After DBRC sets the cursor-active status for the partition in the RECON data sets, and until the copying phase completes and DBRC resets the cursor-active status, you can recover any of the input or output data sets using the Database Recovery utility (DFSURDB0). To restore the output data sets, the Database Recovery utility uses the database change records (type X'50' log records) and applies them to empty output data sets.

Recommendation: Make an image copy of the output data sets as soon as possible after the online reorganization completes. Recovering from this image copy is faster than recovering from the database change records that are logged during the online reorganization. However, you cannot make an image copy while the partition is in cursor-active status.

To recover an output data set before the online reorganization completes, perform the following tasks:

1. Stop the online reorganization by using the `TERMINATE OLREORG` command. If the online reorganization encountered an abend, it is stopped automatically.
2. Issue the `/DBR` or the `UPDATE DB` command for the HALDB partition.
3. Run database change accumulation, as necessary. You can create the JCL by issuing the `GENJCL.CA` command, or you can run the Database Change Accumulation utility (DFSUCUM0) from your own JCL. The purge time for the change accumulation must be equal to the time of the beginning of the online reorganization to represent restoring from the initial empty state of the data set. See "Specifying a Purge Time for the Database Change Accumulation Utility" on page 387.
4. Create the output data set to be recovered, either by using a JCL DD statement or by using Access Method Services, as appropriate.
5. Recover the database changes. You can create the JCL by issuing the `GENJCL.RECOV` command. Alternatively, you can run the Database Recovery utility (DFSURDB0) from your own JCL with the DD statement for DFSUDUMP specified as DUMMY to indicate that there is no image copy from which to restore.
6. Run the Batch Backout utility (DFSBB000), because you might need to back out uncommitted data.
7. After you have recovered, and possibly backed-out, all of the required data sets of the HALDB partition, issue the `/STA DB` or the `UPDATE DB` command for the HALDB partition.
8. Issue the `INITIATE OLREORG` command to resume the online reorganization.

You can also recover an output data set after the online reorganization completes but before an image copy has been made. Follow the same steps as for recovering

an output data set before the online reorganization completes, except the steps for stopping and restarting the online reorganization.

In addition, you can recover an output data set from a point other than the beginning of the online reorganization, such as from a full dump of a DASD volume, using existing procedures if the online reorganization is either completed or terminated.

Specifying a Purge Time for the Database Change Accumulation Utility:

When you run the Database Change Accumulation utility (DFSUCUM0) for one of the output data sets, specify a purge time that is equal to the online reorganization start time.

Specifying this purge time is necessary if change accumulation records (or an input log) that involve the output data set span the time that a online reorganization was started. Specifying the purge time eliminates database change records from before this point in time and is analogous to eliminating database change records from prior to the start time of an image copy.

Specifying a Starting Point for the GENJCL.CA and GENJCL.RECOV

Commands: Even if no image copy exists for the output data sets, the RECON data sets reflect the beginning of the online reorganization as a starting point from which you can perform forward recovery of one of these data sets, even after the online reorganization is complete. Until you make an image copy of an output data set, the GENJCL.CA command treats this starting point as though it were the most recent image copy and causes changes to the output data set to be accumulated from that point. Similarly, the GENJCL.RECOV command prepares recovery of an output data set from this point, even if no physical image copy exists.

Specifying the Active Data Sets for the Database Image Copy Utilities: The database image copy utilities always copy from the currently active data sets that are recorded in the RECON data sets. Regardless of whether the A-through-J or the M-through-V data sets are active, you do not need to change the JCL or control statements for these utilities to specify which set of data sets to use.

On the utility control statement for the Database Image Copy utility (DFSUDMP0), the DDNAME does not need to refer to the currently active data set. Regardless of whether the A-through-J or the M-through-V data sets are active, the utility automatically uses currently active data sets.

Example: Assume that the data set for a second data set group defined in the DBD is to be copied, and that the partition name is PARTNO3. Regardless of which set of data sets is active, you can code a DDNAME of either PARTNO3B or PARTNO3N on the control statement. If the A-through-J data sets are active, whether you specify PARTNO3B or PARTNO3N, the utility copies from PARTNO3B. Likewise, if the M-through-V data sets are active, the utility copies from PARTNO3N.

In the JCL statements for the Database Image Copy utility, you should omit the DD statement that refers to the input data set. Based on whether the A-through-J or the M-through-V data sets are active, the utility dynamically allocates the appropriate data set. A DD statement that refers to a specific data set name can cause the utility job to fail because of a “Data Set Not Found” condition during job-step initiation. This condition occurs if an inactive data set name is coded in the JCL and the data set does not exist.

Activating Sequential Buffering to Improve the Performance of HALDB Online Reorganization

You can use sequential buffering to improve the performance of online reorganization of OSAM databases by including the SBONLINE statement in the IMS.PROCLIB data set member DFSVSMxx.

Using the SBONLINE statement causes IMS to load the sequential buffering modules during initialization so that, whenever you start an online reorganization for an OSAM partition, IMS activates sequential buffering immediately. If you do not include the SBONLINE statement, IMS analyzes the DL/I calls to determine whether sequential buffering is suited for processing the reorganization.

The two forms of the SBONLINE control statement are:

```
SBONLINE  
SBONLINE,MAXSB=nnnnn
```

where *nnnnn* is the maximum amount of storage (in kilobytes) that can be allocated to sequential buffers.

When the maximum amount of storage is reached, IMS stops allocating sequential buffers to online applications (including HALDB Online Reorganization) until these applications release sequential buffer space. If you do not specify the MAXSB= keyword, the maximum amount of storage for sequential buffers is unlimited. For more information about the SBONLINE control statement, see the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

The HALDB Self-Healing Pointer Process

Reorganizations of HALDBs with logical relationships and secondary indexes do not require the execution of utilities to update pointers. Instead, HALDB uses a self-healing pointer process to correct logical relationship and secondary index pointers. This process is implemented by placing a target key and an extended pointer set (EPS) in the secondary index or logically related database and by using an indirect list data set (ILDS) in each partition of PHDAM and PHIDAM databases.

How the Self-Healing Pointer Process Works

The elements of the self-healing pointer process can be seen in Figure 221 on page 389, which shows the interrelationships between the elements prior to a database reorganization.

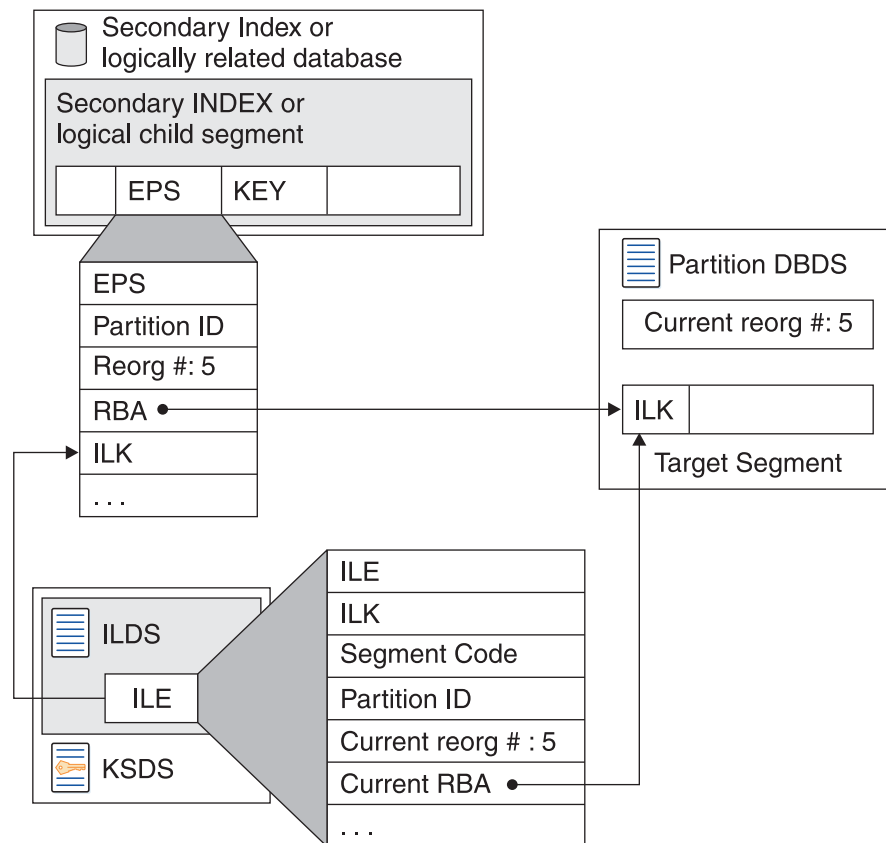


Figure 221. HALDB Pointer Before a Reorganization

Each secondary index entry and each logical child segment contains the key of its target record. For secondary indexes, the key of the target's root segment is included in the prefix. For logical child segments, the concatenated key of the logical parent is included in the segment data.

Each segment in a PHDAM or PHIDAM database has an indirect list key (ILK). The ILK is unique for the segment type across the entire database. It is composed of the relative byte address (RBA), partition ID, and partition reorganization number of the segment when it was first created. The ILK for a segment never changes. It is maintained across reorganizations.

Each secondary index entry or logical child segment has an extended pointer set (EPS). The EPS includes the ILK of its target segment. It also contains the RBA, partition ID, and partition reorganization number for the target segment. These parts of the EPS might not be accurate. That is, they might not reflect the current location of the target segment or the current reorganization number of the target segment's partition. In Figure 221 they are accurate.

The target segment has an indirect list entry (ILE) in the ILDS for a partition. The ILE contains accurate information about the target segment. This includes its current RBA, the correct partition ID, and the current reorganization number for the partition. The key of the ILE is composed of the ILK and the segment code of the target segment.

Reorganizing HALDBs

The reorganization number for a partition is physically stored in the partition's first database data set. This number is initialized by partition initialization or load, and incremented with each reorganization that reloads segments in the partition.

Finding Target Segments

When IMS accesses the target segment from the secondary index entry or logical child segment, it must first determine the partition in which the target resides. It uses the key in the secondary index or logical child to determine the partition. Next it must determine the location in the target partition database data set. It compares the partition ID and reorganization number of the target partition with the partition ID and reorganization number stored in the EPS. If they match, IMS uses the RBA in the EPS to locate the target segment. If they do not match, the RBA in the EPS cannot be used.

When the RBA in the EPS cannot be used, IMS uses the information in the ILE to locate the target segment. The ILE key is found by using the ILK from the EPS and the target's segment code. The ILE is read from the ILDS of the partition determined from the target's key.

Figure 222 on page 391 illustrates a situation in which the RBA in the EPS cannot be used. In the figure, the target partition has been reorganized three times since the EPS was accurate. This has moved the target segment and updated the reorganization number in the partition data set. The EPS still contains a reorganization number of 5, but the reorganization number in the partition data set is now 8. The information in the ILE has been updated by the HD Reorganization Reload utility. IMS uses the ILK from the EPS to find the ILE and uses the RBA in the ILE to find the target segment.

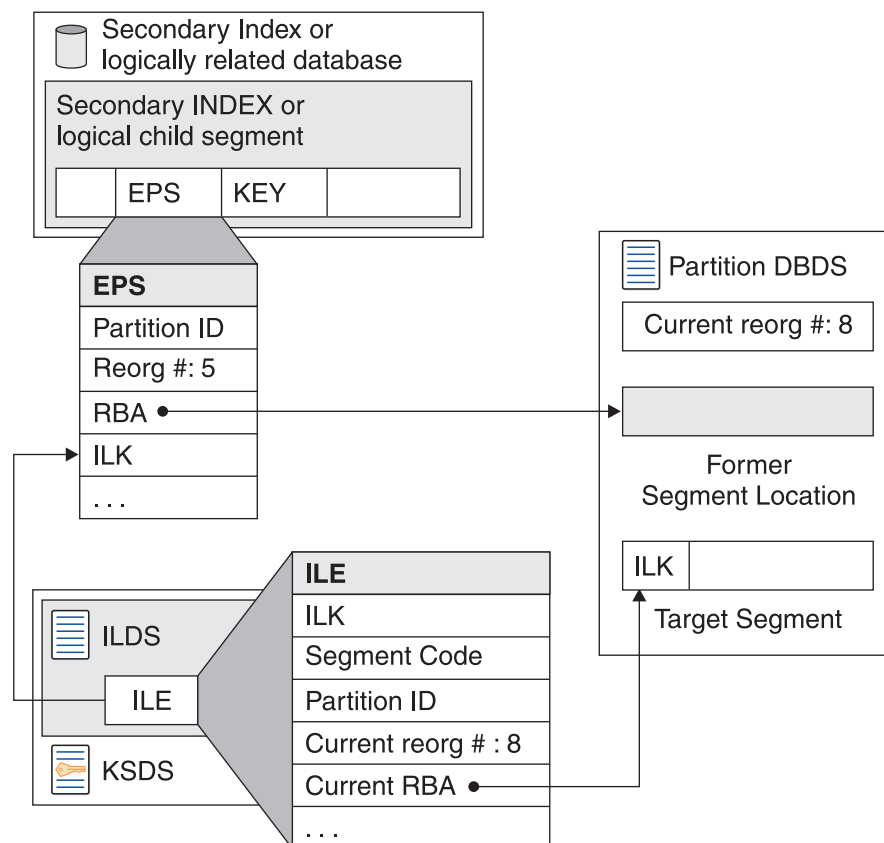


Figure 222. HALDB Pointer After a Reorganization

Even though the retrieval is indirect, often the CI containing the ILE will already be in IMS's buffer pool.

Recommendation: If possible, avoid the indirect process of locating target segments. Instead, get the target segment location from the EPS without reading the ILE. The self-healing process allows IMS to limit the use of ILEs.

Healing Pointers

The self-healing process updates or corrects the information in EPSs. When the ILE is used, the information about the current location of the segment in the ILE is moved to the EPS. This allows IMS to avoid the indirect process if the EPS is used for a later retrieval. This correction to the EPS in the database buffer pool is always done.

Because of locking considerations, the update might not be written to the database on DASD. The buffer containing the entry or segment with the updated EPS is marked as altered if the application program is allowed to update the database. The call must be done with a PCB allowing updates, and the IMS system must have an access intent for the partition that allows updates. If updates are not allowed, the buffer is not marked as altered.

When the application reaches a sync point, it does not write buffers to DASD if they are not marked as altered. If the updated EPS is not written to DASD, the next time it is retrieved from DASD and used to find its target, IMS must use the indirect process. That is, IMS must read the ILE again.

Reorganizing HALDBs

Figure 223 shows the EPS after it has been healed. The RBA points to the current location. The partition ID is correct. The partition reorganization number matches the number stored in the partition database data set.

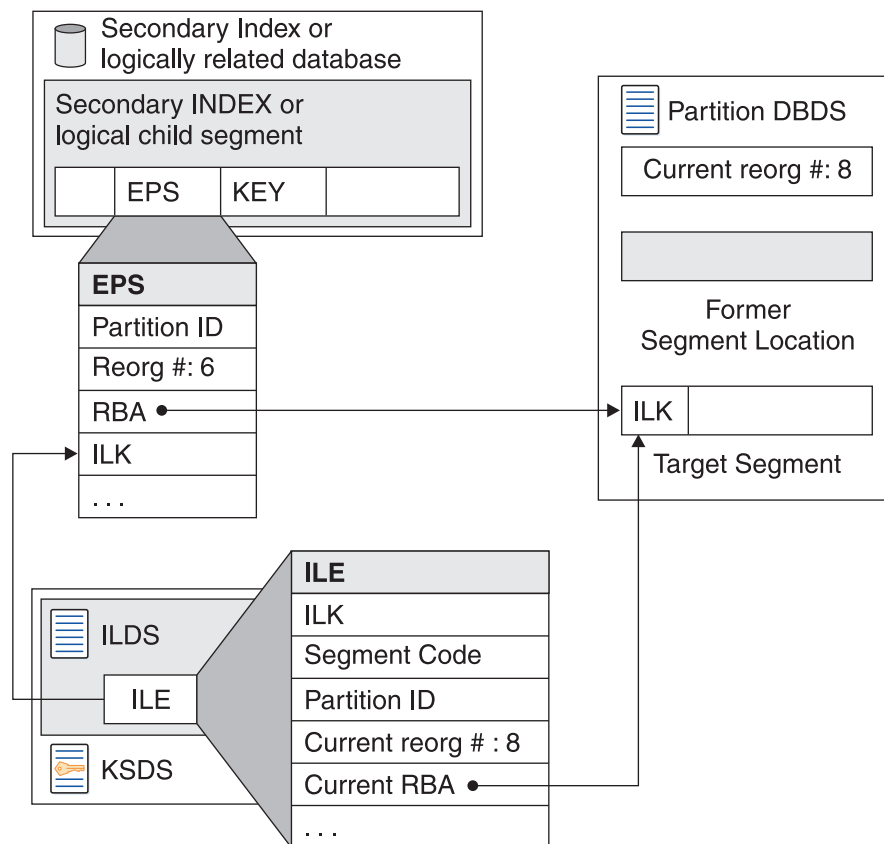


Figure 223. HALDB Pointer After the Self-Healing Process

Performance of the Self-Healing Process

The performance of the self-healing process can be much more efficient than you might anticipate.

Many pointers can be healed with a small number of ILDS reads. This is due to the use of IMS database buffering. ILDSs are database data sets. They use database buffer pools in the same way that other database data sets use them. If a CI is already in its buffer pool, it does not have to be read from DASD.

Each ILE is 50 bytes. You specify the CI sizes for your ILDSs. An 8 KB ILDS CI holds up to 163 ILEs and a 16 KB CI holds up to 327 ILEs, so a single CI can hold many ILEs. After a reorganization, IMS might need to heal many pointers to the reorganized partitions.

When there are frequent uses of the CIs in an ILDS, they tend to remain in their buffer pool. One read of an ILDS CI might be sufficient to heal hundreds of pointers. As with most IMS database tuning, having a large number of buffers for frequently used data sets can be highly beneficial.

Another benefit of the self-healing process is that it does not waste resources healing pointers that are not used. In many secondary indexes, only a small number of entries are actually used. With a non-HALDB database, the entire index is rebuilt

every time the indexed database is reorganized. With HALDB, the index is not rebuilt and only a small number of referenced index entries are updated. HALDB does not use resources to update pointers that are never used.

When an EPS is updated, an application marks buffers as altered only if the application is allowed to make updates. If updates are allowed and a block-level data sharing environment is being used, a block lock is requested for the altered block. Block level data sharing environments exist when the IRLM is used and the share level for the database is either 2 or 3. The block locks are held until the application program commits its unit of work, which could cause a performance problem.

Optimizing Self-Healing Performance: Usually application programs with update authority commit frequently. This is good programming practice. Occasionally, an application program that is allowed to do updates does not actually do them. For example, a program with a PCB specifying PROCOPT=A might only read. In this case, it might not commit frequently. Because it only reads, it never holds many locks. This could change with the implementation of HALDB. If the program runs in a block level data sharing environment and invokes the healing process, it will hold block locks until they are committed. This could cause two problems. First, it might hold the locks for a long time and cause other programs to wait before they can update the blocks. Second, it could hold many locks. This could cause a storage shortage in the IRLM or a lock structure.

If you have a program that holds locks for a long time or that holds many locks when performing the self-healing pointer process, you have four options:

- If the application program does not make updates, use PROCOPT=G.
- Have your program commit frequently.
- Invoke the pointer healing process before you run application programs that use PROCOPT=A, but do not do any updates. Run another program or utility before this type of application program. The HALDB Conversion and Maintenance Aid tool supplies a pointer healing utility.
- Rebuild secondary indexes with an index builder, such as the IMS Index Builder for z/OS. The IMS Index Builder for z/OS creates EPSs with accurate RBAs.

This scenario is not common. Most users can let the pointer healing process occur without taking any special precautions.

Recommendation: Do not rebuild your secondary indexes after a reorganization. Let the self-healing process of HALDB correct the pointers. This shortens the outage for reorganizations and tends to minimize the use of resources.

Related Reading:

- For more information about the IMS High Availability Large Database Conversion and Maintenance Aid, see the *IMS High Availability Large Database Conversion and Maintenance Aid for z/OS, User's Guide*.
- For more information about the IMS Index Builder, see the *IMS Index Builder for z/OS User's Guide*.

Changing DL/I Access Methods

When you originally chose a DL/I access method (or type of database), you chose it based on such things as:

- The type of processing you needed to do (sequential, direct, or both)
- The volatility of your data

If the characteristics of your applications have changed over a period of time, performance might be improved by changing to another DL/I access method. Chapter 6, “Choosing Full-Function Database Types,” on page 59 describes which type of DL/I access method to choose given your application’s characteristics.

Assuming that you have decided to change access methods, this topic tells you:

- Given your existing DL/I access method, what things you need to change to convert to a different DL/I access method
- How to do the conversion

The reorganization utilities described earlier in this chapter can be used to change DL/I access methods among the HISAM, HDAM, and HIDAM access methods. One exception to this is that HDAM cannot be changed to HISAM or HIDAM unless HDAM database physical records are in root key sequence. This exception exists because HISAM and HIDAM databases must be loaded with database records in root key sequence. When the HD Reorganization Unload utility unloads an HDAM database, it unloads it using GN calls. GN calls against an HDAM database unload the database records in the physical sequence in which they were stored by the randomizing module. This will not be root key sequence unless you used a sequential randomizing module (one that put the database records into the database in physical root key sequence).

Related Reading: The procedures in this topic require you to reassess different aspects of your databases. See the following related readings for information to help you make the reassessments:

- For a description of free space and how it is specified, see “Specifying Free Space (HDAM, PHDAM, HIDAM, and PHIDAM Only)” on page 245.
- For a description of types of pointers and how to specify them, see “Types of Pointers You Can Specify” on page 85.
- For information about what to consider in choosing a logical record length and how logical record lengths are specified, see “Choosing a Logical Record Length for HD Databases” on page 252.
- For information about what to consider in choosing a CI or block size and how CI and block size are specified, see “Determining the Size of CIs and Blocks” on page 252.
- For information about what to consider in choosing buffer number and size and how buffers are specified, see “Buffer Numbers” on page 255.
- For information about how to calculate database size, see “Estimating the Minimum Size of the Database” on page 315.
- For information about choosing HDAM or PHDAM options, see “Choosing HDAM or PHDAM Options” on page 248.
- For information about choosing and specifying a randomizing module, see “Determining Which Randomizing Module to Use (HDAM and PHDAM Only)” on page 247.

Changing the DL/I Access Method From HISAM to HIDAM

You need the following before changing your DL/I access method from HISAM to HIDAM:

- Determine whether you are going to set aside free space in the HIDAM database. (Free space is space into which database records are not loaded when the database is initially loaded.)
Unlike HISAM, in a HIDAM database you can set aside periodic blocks or CIs of free space or a percentage of free space in each block or CI (in the ESDS or OSAM data set). This free space can then be used for inserting database records or segments into the database after initial load.
- Determine what type of pointers you are going to use in the database. Unlike HISAM, HIDAM uses direct-address pointers to point from one segment in the database to the next.
- Reassess your choice of logical record size. A logical record in HISAM can only contain segments from the same database record. In HIDAM, a logical record can contain segments from more than one database record.
- Reassess your choice of CI or block size. In HISAM, your choice of CI or block size should have been some multiple of the average size of a database record. In HIDAM, the size should be chosen because of the characteristics of the device and the type of processing you plan to do.
- Reassess your choice of database buffer sizes and the number of buffers you have allocated. If you have changed your CI or block size, you need to allocate buffers for the new size.
- Recalculate database space. You need to do this because the changes you are making will result in different requirements for database space.

Once you have determined what changes you need to make, you are ready to change your DL/I access method from HISAM to HIDAM. To do this:

1. Unload your database using the existing DBD and the HD Reorganization Unload utility.
2. Code a new DBD that reflects the changes you need to make. You must also code a DBD for the HIDAM index.
3. If you need to make change that are not specified in the DBD (such as changing database buffer sizes or the amount of space allocated for the database), make these changes.
4. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
5. Reload the database using the new DBD and the HD Reorganization Reload utility. Remember to make an image copy of your database as soon as it is reloaded.

If you are using logical relationships or secondary indexes, you will need to run additional utilities immediately before and after reloading your database. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing the DL/I Access Method From HISAM to HDAM

You need to do the following before changing your DL/I access method from HISAM to HDAM:

Changing DL/I Access Methods

- Determine what type of pointers you are going to use in the database. Unlike HISAM, HDAM uses direct-address pointers to point from one segment in the database to the next.
- Determine which randomizing module you are going to use. Unlike HISAM, HDAM uses a randomizing module. The randomizing module generates information that determines where a database record will be stored.
- Determine which HDAM options you are going to use. Unlike HISAM, an HDAM database is divided into two parts: a root addressable area and an overflow area. The root addressable area contains all root segments and is the primary storage area for dependent segments in a database record. The overflow area is for storage of dependent segments that do not fit in the root addressable area. The HDAM options here are the ones that pertain to choices you make about the root addressable area. These are:
 - The maximum number of bytes of a database record to be put in the root addressable area when segments in the database record are inserted consecutively (without intervening processing operations).
 - The number of blocks or CIs in the root addressable area.
 - The number of RAPS (root anchor points) in a block or CI in the root addressable area. (A RAP is a field that points to a root segment.)
- Reassess your choice of logical record sizes. A logical record in HISAM can only contain segments from the same database record. In HDAM, a logical record can contain segments from more than one database record. In addition, HDAM logical records contain RAPS and two space management fields (FSEs and FSEAPs).
- Reassess your choice of CI or block size. In HISAM, your choice of CI or block size should have been some multiple of the average size of a database record. In HDAM, the size should be chosen because of the characteristics of the device and the type of processing you plan to do.
- Reassess your choice of database buffer sizes and the number of buffers you have allocated. If you have changed your CI or block size, you need to allocate buffers for the new size.
- Recalculate database space. You need to do this because the changes you are making will result in different requirements for database space.

Once you have determined what changes you need to make, you are ready to change your DL/I access method from HISAM to HDAM. To do this:

1. Unload your database, using the existing DBD and the HD Reorganization Unload utility.
2. Code a new DBD that reflects the changes you need to make.
3. If you need to make changes that are not specified in the DBD (such as changing database buffer sizes or the amount of space allocated for the database), make these changes. HDAM only requires one data set, whereas HISAM requires two.
4. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
5. Reload the database using the new DBD and the HD Reorganization Reload utility. Make an image copy of your database as soon as it is reloaded.

If you are using logical relationships or secondary indexes, you will need to run additional utilities before reloading your database. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing the DL/I Access Method From HIDAM to HISAM

You need to do the following before changing your DL/I access method from HIDAM to HISAM:

- Reassess your choice of logical record size. A logical record in HISAM can only contain segments from the same database record. In HIDAM, a logical record can contain segments from more than one database record.
- Reassess your choice of CI or block size. In HIDAM, your choice of CI or block size should be based on the characteristics of the device and the type of processing you plan to do. In HISAM, the size should be some multiple of the average size of a database record.
- Reassess your choice of database buffer sizes and the number of buffers you have allocated. If you have changed your CI or block size, you need to allocate buffers for the new size.
- Recalculate database space. You need to do this because the changes you are making will result in different requirements for database space.

Once you have determined what changes you need to make, you are ready to change your DL/I access method from HIDAM to HISAM. To do this:

1. Unload your database using the existing DBD and the HD Reorganization Unload utility.
2. Code a new DBD that reflects the changes you need to make. You will not be specifying direct-address pointers or free space in the DBD, because HISAM, unlike HIDAM, does not allow use of these. Also, HISAM has only one DBD whereas HIDAM had two.
3. If you need to make changes that are not specified in the DBD (such as changing database buffer sizes or the amount of space allocated for the database), make these changes.
4. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
5. Reload the database using the new DBD and the HD Reorganization Reload utility. Remember to make an image copy of your database as soon as it is reloaded.

If you are using logical relationships or secondary indexes, run additional utilities right before and after reloading your database. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing the DL/I Access Method From HIDAM to HDAM

You need to do the following before changing your DL/I access method from HIDAM to HDAM:

- Reassess your choice of direct-address pointers. Although both HIDAM and HDAM use direct-address pointers, you might need to change the type of direct-address pointer used:
 - Because of the changing needs of your applications.
 - Because pointers are partly chosen based on the type of database you are using. For example, if you used physical twin backward pointers on root segments in your HIDAM database to get fast sequential processing of roots, they will not have any use in an HDAM database. See Chapter 6, “Choosing Full-Function Database Types,” on page 59 under “Types of Pointers You Can Specify” for a description of types of pointers, their uses, and how to specify them.

Changing DL/I Access Methods

- Determine which randomizing module you are going to use. Unlike HIDAM, HDAM uses a randomizing module. The randomizing module generates information that determines where a database record is to be stored.
- Determine which HDAM options you are going to use. Unlike HIDAM, an HDAM database does not have a separate index database. Instead the database is divided into two parts: a root addressable area and an overflow area. The root addressable area contains all root segments and is the primary storage area for dependent segments in a database record. The overflow area is for storage of dependent segments that do not fit in the root addressable area. The HDAM options here are the ones that pertain to choices you make about the root addressable area. These are:
 - The maximum number of bytes of a database record to be put in the root addressable area when segments in the database record are inserted consecutively (without intervening processing operations).
 - The number of blocks or CIs in the root addressable area.
 - The number of RAPs in a block or CI in the root addressable area.
- Reassess your choice of logical record size.
- Reassess your choice of CI or block size.
- Reassess your choice of database buffer sizes and the number of buffers you have allocated. If you have changed your CI or block size, you need to allocate buffers for the new size.
- Recalculate database space. You need to do this because the changes you are making will result in different requirements for database space.

After you have determined what changes you need to make, you are ready to change your DL/I access method from HIDAM to HDAM. To do this:

1. Unload your database using the existing DBD and the HD Reorganization Unload utility.
2. Code a new DBD that reflects the changes you need to make. You probably will not be specifying free space, but you will be specifying HDAM options. Note also that you'll need only one DBD for HDAM, whereas HIDAM required two DBDs.
3. If you need to make changes that are not specified in the DBD (such as changing database buffer sizes or the amount of space allocated for the database), make these changes. HDAM only requires one data set, whereas HIDAM requires two.
4. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
5. Reload the database using the new DBD and the HD Reorganization Reload utility. Remember to make an image copy of your database as soon as it is reloaded.

If you are using logical relationships or secondary indexes, you will need to run additional utilities right before and after reloading your database. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing the DL/I Access Method From HDAM to HISAM

You need to do the following before changing your DL/I access method from HDAM to HISAM:

- Reassess your choice of logical record size. A logical record in HISAM can only contain segments from the same database record. In HISAM, a logical record can contain segments from more than one database record.
- Reassess your choice of CI or block size. In HDAM, your choice of CI or block size should be based on the characteristics of the device and the type of processing you plan to do. In HISAM, the size should be some multiple of the average size of a database record.
- Reassess your choice of database buffer sizes and the number of buffers you have allocated. If you have changed your CI or block size, you need to allocate buffers for the new size.
- Recalculate database space. You need to recalculate database space because the changes you are making will result in different requirements for database space.

After you have determined what changes you need to make, you are ready to change your DL/I access method from HDAM to HISAM. Remember you must write your own unload and reload programs unless database records in the HDAM database are in physical root key sequence. In writing your own load program, if your HDAM database uses logical relationships, you must preserve information in the delete byte (for example, a segment that is logically deleted in the database might not be physically deleted).

To change from HDAM to HISAM:

1. Unload your database using the existing DBD and one of the following:
 - Your unload program
 - The HD Reorganization Unload utility if database records are in physical root key sequence
2. Code a new DBD that reflects the changes you need to make. You will not be specifying direct-address pointers or HDAM options.
3. If you need to make changes that are not specified in the DBD (such as changing database buffer sizes or the amount of space allocated for the database), make these changes. HDAM only requires one data set, whereas HISAM requires two.
4. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
5. Reload the database using the new DBD and:
 - Your load program, or
 - The HD Reorganization Reload utility if database records are in physical root key sequences

Remember to make an image copy of your database as soon as it is reloaded.

If you are using logical relationships or secondary indexes, you will need to run additional utilities right before and after reloading your database. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing the DL/I Access Method From HDAM to HIDAM

You need to make the following changes before changing your DL/I access method from HDAM to HIDAM:

- Determine whether you are going to set aside free space in the HIDAM database. (Free space is space into which database records are *not* loaded)

Changing DL/I Access Methods

when the database is initially loaded.) In a HIDAM database, you can set aside periodic blocks or CIs of free space or a percentage of free space in each block or CI (in the ESDS or OSAM data set). This free space can then be used for inserting database records or segments into the database after initial load. In an HDAM database, you generally get the free space you need by careful choice of HDAM options.

- Reassess your choice of direct-address pointers. Although both HIDAM and HDAM use direct-address pointers, you might need to change the type of direct-address pointer used:
 - Because of the changing needs of your applications.
 - Because pointers are partly chosen based on the type of database you are using. For example, you can choose to use physical twin forward and backward pointers on root segments in your HIDAM database to get fast sequential processing of roots.
- Reassess your choice of logical record size.
- Reassess your choice of CI or block size.
- Reassess your choice of database buffer sizes and the number of buffers you have allocated. If you have changed your CI or block size, you need to allocate buffers for the new size.
- Recalculate database space. You need to recalculate database space because the changes you are making will result in different requirements for database space.

Once you have determined what changes you need to make, you are ready to change your DL/I access method from HDAM to HIDAM. Remember you must write your own unload and reload programs unless database records in the HDAM database are in physical root key sequence. In writing your own load program, if your HDAM database uses logical relationships, you must preserve information in the delete byte (for example, a segment that is logically deleted in the database might not be physically deleted).

To change from HDAM to HIDAM:

1. Unload your database using the existing DBD and one of the following:
 - Your unload program
 - The HD Reorganization Unload utility if database records are in physical root key sequence
2. Code a new DBD that reflects the changes you need to make. You must also code a DBD for the HIDAM index. You will not be specifying HDAM options but you probably will be specifying free space.
3. If you need to make changes that are not specified in the DBD (such as changing database buffer sizes or the amount of space allocated for the database), make these changes. HDAM only requires one data set, whereas HIDAM requires two.
4. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
5. Reload the database using the new DBD and one of the following:
 - Your load program
 - The HD Reorganization Reload utility if database records are in physical root key sequence.

Remember to make an image copy of your database as soon as it is reloaded.

If you are using logical relationships or secondary indexes, you will need to run additional utilities before reloading your database. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing the DL/I Access Method From HDAM to PHDAM and HIDAM to PHIDAM

For a logical view of HDAM and HIDAM databases before and after changing to PHDAM and PHIDAM respectively, see Figure 224.

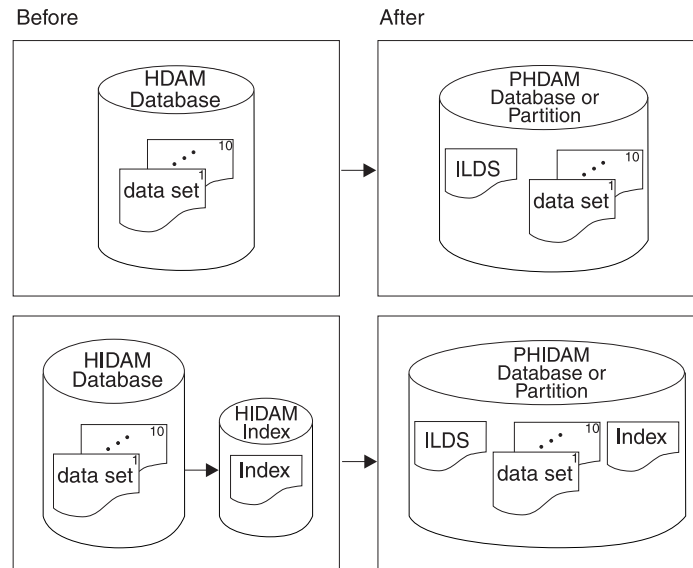


Figure 224. HDAM and HIDAM Databases Before and After Changing to PHDAM and PHIDAM

Requirement: You must concurrently migrate *all* databases that are logically related. All secondary indexes that point to these logically related databases must be migrated at the same time the databases they point to are migrated.

Because non-keyed PHDAM root segments are not supported, you cannot migrate an HDAM database with non-keyed roots to HALDB.

There are two methods for changing a HDAM or HIDAM database to PHDAM or PHIDAM. The first method keeps the same database name. The second method changes the name of the physical database and uses a logical database with the old database name.

Secondary Index Considerations

Migration of secondary index databases with non-unique keys requires separate JCL steps to sort and merge the unload records to create new /SX values prior to inputting them into the HD Reorganization Reload utility. User data is lost.

Related Reading: For more information about the HD Reorganization Reload Utility, see *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

Secondary indexes targeting HDAM or HIDAM databases that are changing to PHDAM or PHIDAM must be changed to PSINDEXs. The steps for doing this are

Changing DL/I Access Methods

the same as the steps for unload and reload. Run the HD Reorganization Unload and Reload utilities against the secondary index. The user data is preserved in the secondary index.

Determining the Database Name

When migrating a database to a HALDB format, you must decide whether to change the database name.

If the new database is to have the same name as the old database:

1. Unload the old database with the migrate option before changing RECON or DBDLIB.
2. Create a RECON list before deleting the records for the database.
3. Remove the information from the old database RECON and DBDLIB.
4. Delete all MDA members
5. Define the HALDB by using DBDGEN, ACBGEN, and either the HALDB Partition Definition utility or the DBRC commands INIT.DB and INIT.PART.

After the definitions are complete, load the new database.

If the new database is to have a different name from the old database:

1. Create a RECON list before deleting the records for the database. The old information is retained in RECON as long as necessary.
2. Unload the old database.
3. Remove the DBD from DBDLIB and ACBLIB.
4. Delete all MDA members that refer to the old database.
5. Perform a DBDGEN on the old database name as a logical database with the source being the new HALDB.
6. Define the HALDB by using DBDGEN, ACBGEN, and either the HALDB Partition Definition utility or the DBRC commands INIT.DB and INIT.PART.

Performing the Migration Unload

Use the HD Reorganization Unload utility (DFSURGU0) to unload the database. Specify the migrate option to unload the database.

Performing the Migration Reload

To reload the database:

1. Run the DBDGEN utility for the new HALDB DBD.
2. Use either the HALDB Partition Definition utility or the DBRC commands INIT.DB and INIT.PART to define the new HALDB partitions.
3. Run the HALDB Partition Data Set Initialization utility (DFSUPNT0) or the Database Prereorganization utility (DFSURPR0).
4. Reload the database by using HD Reorganization Reload utility (DFSURGL0).
5. Make an image copy of all partitions.

The user data is preserved in the secondary index.

Changing the DL/I Access Method From PHDAM and PHIDAM to HDAM and HIDAM

The process of restoring HDAM or HIDAM databases that were migrated to PHDAM or PHIDAM is known as fallback. Fallback supports the following types of logical relationships:

- Unidirectional HALDB to current unidirectional database

- Physically paired HALDBs to current physically paired databases

The order of physical twin segments is maintained when a fallback from HALDBs occurs. This includes segments that are non-keyed and that have a non-unique key.

Primary indexes are recreated, not unloaded. Secondary indexes are recreated by the reload utility process. User data is not preserved.

Requirements: The requirements for a fallback include:

- You must perform a concurrent fallback of all databases that are logically related.
- You must have prefix resolution and prefix update utilities, if logical children or secondary indexes are present when a database falls back from HALDB.
- Before you use the database, but after you reload and perform any prefix resolution or prefix update, take an image copy of all data sets, including the prime index, by using one of the image copy utilities. The image copy utilities invoke DBRC to validate input and record results, which ensures that you have a backup copy of the database that can serve as an effective point of recovery in case of failure.
- You must complete the fallback of all related databases and secondary index databases before any database can be used.

Restriction: You cannot perform a fallback on physically paired HALDBs to current virtually paired databases and preserve the logical sequence of the virtual logical child. The step to accomplish this conversion are:

1. Perform a fallback on current physically paired databases.
2. Reorganize the current database.
3. Change the logical relationship to virtually paired databases.

The steps in the fallback process are:

1. Unload all related databases using the HD Reorganization Unload utility (DFSURGU0) with the FALLBACK option. DFSURGU0 locates the paired logical children and saves information needed for fallback in the prefix of the output data. The prefix created by DFSURGU0 contains the information required to create the new segment prefix when the data is reloaded.
2. Perform DBDGENs to define the current format databases.
3. Re-register all databases to be controlled by DBRC. If keeping the same database name, first use the HALDB partition Definition utility to delete the HALDB.
4. Perform the prereorganization step with the databases listed as DBR.
5. Reload all the unloaded databases using the new definitions.

Logical children have some special considerations. There are three cases to consider: unidirectional, virtually paired, and physically paired databases. Current DL/I offers an option to not store the logical parent's concatenated key in the logical child (virtual key storage option); in normal retrieval the key is built and the user application always sees the concatenated key in the data. For all logical children unloaded, you must drop the logical parent's concatenated key if the virtual key storage option is chosen. The unloaded segments are reloaded as real segments that are part of a physically paired relationship. This type of unload, dropping the logical parent's concatenated key, only occurs when DFSURGU0 performs a fallback unload.

Changing HALDB Partition Definitions

HALDB partitions can require changing when they become too large, too small, empty, obsolete, and so forth. In most cases, HALDB partitions can be added, deleted, enabled, or disabled without requiring the whole database to be unavailable. In other cases, you must take the entire HALDB offline to make changes.

This topic discusses the following subjects:

- “Making Changes to a Single Partition”
- “Making Changes that Affect All of the Partitions in a HALDB”
- “Updating Partition Structures When Changing Partition Definitions” on page 405
- “Changing Partition Boundaries” on page 406

Making Changes to a Single Partition

Certain characteristics of HALDBs are specific to each partition in the HALDB. To change these characteristics, you do not need to take the entire HALDB offline; you only need to take the partition in which you are making changes offline. Before making changes to a partition, issue a /DBRECOVERY command or an UPDATE DB STOP (ACCESS) command against the partition.

You can change the following characteristics of a single partition:

- DSN prefix
- Randomizing module name
- Number of root anchor points (RAPs)
- Bytes parameter
- OSAM block size
- VSAM CI size

For example, suppose a PHDAM partition has many roots that randomize to the same root anchor point. This causes lock contention problems that negatively impact performance. To remedy this problem, you can increase the number of RAPs in the partition.

The following steps describe how to change the number of RAPs in a partition:

1. Issue the /DBRECOVERY command to take the partition offline.
2. Unload the data from the partition.
3. Use the Partition Definition utility or DBRC commands to change the number of RAPs in the partition.
4. Reload the partition.
5. Take an image copy of the data sets for the partition.
6. Issue the /START DB command or the UPDATE DB command to make the partition available again.

Making Changes that Affect All of the Partitions in a HALDB

Certain characteristics of HALDB partitions are common to all partitions in a HALDB. Before you can change these characteristics, you must take all of the partitions in the HALDB offline by issuing the /DBRECOVERY command.

You can change the following characteristics of HALDB partitions only after taking all partitions offline:

- DBD definition

- The HALDB Partition Selection exit routine (DFSPSE00)
- Share level
- Nonrecoverable attribute status
- RSR GSG name or tracking level

For example, suppose a HALDB has an existing HALDB Partition Selection exit routine that needs to be replaced with a HALDB Partition Selection exit routine that selects partitions based on a new algorithm. This change requires the entire HALDB to be offline.

The steps below describe how to change a HALDB Partition Selection exit routine:

1. Issue the /DBRECOVERY command to take the HALDB offline.
2. Unload the data from the HALDB using the existing HALDB Partition Selection exit routine.
3. Use the Partition Definition utility or DBRC commands to change the HALDB Partition Selection exit routine.
4. Reload the data from the HALDB using the new HALDB Partition Selection exit routine.
5. Run Image Copy for the data sets for all partitions in the HALDB.
6. Issue the /START command or the UPDATE DB command to make the HALDB available again.

Updating Partition Structures When Changing Partition Definitions

The online IMS system creates an internal *partition structure* for each HALDB partition you define. Any time you change HALDB partition definitions in DBRC's RECON data set, IMS must rebuild the partition structure of all partitions affected by the change before the partitions can be used.

To detect when IMS needs to rebuild a partition structure, IMS tracks changes to HALDBs using a *change version number*. Each time you change the partition definitions recorded in the RECON data set, DBRC increases the change version number for the HALDB master database by one. When the online IMS system detects an increase in the change version number for a HALDB, IMS rebuilds the partition structure of each partition affected by the change.

Online change is not used for changing HALDB partition definitions or rebuilding partition structures. IMS recognizes any increases in HALDB change version numbers and dynamically reflects the new definitions in the online IMS system.

If you are using XRF, the alternate IMS system sees the dynamic change and automatically updates the definitions in the alternate system, requiring no action from you.

Events That Trigger a Structure Rebuild: IMS checks the HALDB change version numbers to detect changes to partition definitions and rebuild partition structures, if necessary, in the following events:

- When a partition is authorized for use. This would detect a change in an existing partition or in a partition which was not already authorized. This occurs commonly when a partition is taken offline, changes are made to its definition in the RECON data set, and the partition is made available again. The first use of the updated partition triggers partition structure rebuild.

- When an invalid key is detected by partition selection or by a Partition Selection exit routine. This can occur, for example, when a new partition is added beyond the high key of the last partition and all existing partitions are already authorized. In this case, IMS partition selection or the Partition Selection exit routine detects the new partition. After the new partition is detected, IMS performs partition structure rebuild automatically.
- When a `/START DB HALDB_Master OPEN` or `UPDATE DB NAME(HALDB_Master) OPTION(OPEN)` command is issued. For example, if a new partition has been added beyond the high key of the last partition and all existing partitions are already authorized, these commands will initiate a partition structure rebuild. For more details on the `/START DB` command and the `UPDATE DB` command, see *IMS Version 9: Command Reference*.

Additional Considerations Regarding Partition Structure Rebuild: When making changes to HALDB partition definitions, consider the following points:

- If you use a HALDB Partition Selection exit routine, you must issue the `/DBRECOVERY` or the `UPDATE` command and then the `/START` or the `UPDATE` command after making any changes to a partition definition. When the HALDB Partition Selection exit routine selects a HALDB partition, IMS is not aware of HALDB partition boundaries until a partition structure rebuild is complete. For more information on when IMS rebuilds partition structures, see “Events That Trigger a Structure Rebuild” on page 405.
- If you are using a HALDB Partition Selection exit routine and IMS notifies you of a structure modification, you might need the exit routine to select partitions correctly based on the current partition structure.
- Issuing a `/START DB` command with the `OPEN` keyword might fail after the definition of a partition structure has been changed. This is because the online partition structure needs to be rebuilt. To invoke a structure rebuild, an application program that uses the partition must be run or the type-1 command `/START DB HALDB_Master OPEN` must be issued.
- Newly added partitions will not be known by the online IMS system until partition structure rebuild has been invoked and the new structure has been created.

Changing Partition Boundaries

Changes to partition boundaries can affect one or more partitions. If a change to one partition causes records to be moved to or from any other partitions, the change also affects those other partitions. If you use high keys for partition selection, IMS automatically sets the initialization-required flag for the partitions affected by a boundary change. If you use a HALDB Partition Selection exit routine, you are responsible for flagging the partitions that are affected by a boundary change as requiring initialization. In either case, before you change the partition boundaries, you must issue the `/DBR` command against all of the partitions that will be affected by the change.

For example, suppose a HALDB partition named PART200 has a key range from 101 up to a high key of 200 (KEY200). PART200 needs to be split into two HALDB partitions so that a new partition named PART150 is added between another partition, PART100, and PART200. PART150 will have a key range from 101 up to a high key of 150 (KEY150), a key range that used to be included in PART200.

To create a new partition named PART150 by splitting PART200:

1. Take partition PART200 offline by issuing the `/DBRECOVERY` command.
2. Unload the data from PART200.
3. Define the new partition named PART150 with a high key of KEY150.

4. Physically allocate the necessary data sets for PART150.
5. Initialize partitions PART150 and PART200 by running either of the following utilities:
 - HALDB Partition Data Set Initialization utility (DFSUPNT0)
 - Database Prereorganization utility (DFSURPR0)
6. Reload the PART200 data into the data sets that were allocated in step 4. IMS uses the new HALDB partition definitions in DBRC to load the data into PART150 and PART200. With the new definitions, the first root key higher than KEY150 is the first record loaded into PART200.
7. Run an Image Copy utility for the data sets for both PART150 and PART200. DBRC does not allow updates to the data after the image copy required flag is set and before the image copies have been recorded in the RECON data sets.
8. Issue the /START command to PART200 to make it available again.

In this example, online IMS systems do not know of PART150 until one of the following events occur:

- A /START DB *HALDB_Master* OPEN command is issued
- An UPDATE DB *HALDB_Master* OPEN command is issued
- A DL/I call causes an authorization call to DBRC for PART200. The first DL/I call goes through HALDB partition selection again to properly select and authorize either PART150 or PART200.

Procedure for Changing to DEDBs

If your database requires logical relationships, a secondary index, or fixed-length segments, DEDBs cannot be used.

You need to do the following before changing your database to DEDBs:

- Determine whether or not your application programs can tolerate the FH (data unavailable) status code.
- Determine whether or not your database can tolerate a randomizing routine (might not be a problem when changing from HDAM).
- Recalculate database space, particularly when using DEDB features such as partitioning and data set replication.
- Determine which pointers are available to use.

To change to DEDBs:

1. Unload your database using the existing DBD and one of the following:
 - Your unload program
 - The HD Reorganization Unload utility if database records are in physical root key sequence
2. Code a new DBD for the DEDBs.
3. Execute the DBD generation.
4. For non-VSAM data sets, delete the old database space and define the new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
5. Run the DEDB initialization utility (DBFUMIN0).
6. Run the user DEDB load program.

Changing the Hierarchic Structure

There are two types of tuning changes you might need to make that involve changes to the structure of your database record. The first is changing the hierarchic sequence of segment types in your database record to improve performance. The second is combining segments to maximize the use of space.

Changes involving adding and deleting segments in the hierarchy are covered in Chapter 16, “Modifying Databases,” on page 429.

Changing the Sequence of Segment Types

In general, performance is best if frequently used dependent segments are close to the root segment and infrequently used dependent segments are toward the end of the database record. This arrangement maximizes performance because all types of databases (except HSAM) have direct (therefore, fast) access to root segments. But, after the root is located, dependent segments are found by one of the following:

- Searching sequentially through the database record (HSAM and HISAM)
- Following pointers from the root segments to a dependent path and then searching through twin chains until the correct segment is reached (HDAM, HIDAM, PHDAM, and PHIDAM).

One way to determine whether the order of dependent segment types in your hierarchy is an efficient one is to examine the IWAITS/CALL field on the DL/I Call Summary report.

Related Reading: For detailed information on the DL/I Call Summary report, see *IMS Version 9: Utilities Reference: Database and Transaction Manager*.

The IWAITS/CALL field tells you, by DL/I call against a specific segment, the average number of times a segment had to wait for I/O operations to finish before the segment could be processed. A high number (and high, of course, is relative to the application) indicates that multiple I/O operations were required to process the segment.

If the database does not need to be reorganized, the high number can mean this is a frequently used segment type placed too far from the beginning of the database record. If you determine this is the situation, you can change placement of the segment type. The change can increase the value in the IWAITS/CALL field for other segments.

To change the placement of a segment type, you must write a program to unload segments from the database in the new hierarchic sequence. (The reorganization utilities cannot be used to make such a change.) Then you need to load the segments into a new database. Again, you must write a program to reload.

Combining Segments

The second type of change you might need to make in the structure of your database record is combining segment types to maximize use of space. For example, having two segment types, a dependent segment for college classes with a dependent segment for instructors who teach the classes, is an inefficient use of space if typically only one or two instructors teach a class. Rather than having a separate instructor segment, you can combine the two segment types, thereby saving space.

Combining segments also requires that you write an unload and reload program. (The reorganization utilities cannot be used to make such a change.)

Procedure for Changing the Hierarchic Structure

To change the hierarchic structure, you need to:

1. Determine whether the change you are making will affect the code in any application programs. If so, make sure the code gets changed.
2. Unload your database using your unload program and the existing DBD.
3. Code a new DBD.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
6. Reload your database using your load program and the new DBD. Remember to make an image copy of your database as soon as it is reloaded.
7. If your database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after reloading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing Direct-Access Storage Devices

Several situations might warrant tuning your database by changing DASDs (direct-access storage devices). First, when application requirements change, you might require a faster or slower device. Second, you might want to take advantage of new devices offering better performance. Finally, you might need to change devices to get database data sets on two different devices, so as to minimize contention for device use.

You can change your database (or part of it) from one device to another using the reorganization utilities. To change direct-access storage devices:

1. Unload your database using the existing DBD and the appropriate unload utility.
2. Recalculate CI or block size to maximize use of track space on the new device. Information on calculating CI or block size is contained in Chapter 9, "Designing Full-Function Databases," on page 245 under "Determining the Size of CIs and Blocks".
3. Code a new DBD.
4. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
5. Reload your database, using the new DBD and the appropriate reload utility. Remember to make an image copy of your database as soon as it is reloaded.
6. If your database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after reloading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Tuning OSAM Sequential Buffering

If you are using OSAM Sequential Buffering, you can do two things to help ensure that it processes your databases efficiently:

- Keep your databases well organized; that is, the logical (database record) sequence is nearly the same as the physical (DASD block) sequence.
- Select the right number of SB buffer sets (Tuning of SB buffers is discussed “OSAM Sequential Buffering” on page 414).

Well-Organized Database

Well-organized databases are by far the most important of these two factors. When the databases SB processes are well organized, you note elapsed time improvements. This is because your programs process IMS database segments and records, and they do not process DASD blocks directly. Processing a well-organized database in logical-record sequence results in an I/O reference pattern that accesses most DASD blocks in physical sequence. SB can take advantage of these sequential I/O patterns by issuing many sequential reads. Extensive use of sequential reads considerably reduces the elapsed time for your job.

Badly-Organized Database

Processing a badly-organized database in logical-record sequence typically results in an I/O reference pattern that accesses many DASD blocks in a random sequence. This happens because many segments were stored in randomly scattered blocks after the database was loaded or reorganized. When your database is accessed in a predominantly random pattern, most I/O operations issued by the SB buffer handler are random reads. SB is not able to issue many sequential reads, and the elapsed time for your job is not considerably reduced.

You can use the SB buffering statistics in the optional //DFSSTAT reports to see if your database is well-organized. Your database is likely to be badly organized if a large percentage of the blocks were read with random reads during sequential processing. You can monitor this percentage over a period of time to see if it increases as the database ages.

Related Reading: For details on //DFSSTAT reports, see *IMS Version 9: Utilities Reference: System*.

Ensuring a Well-Organized Database

You can ensure your databases are reasonably well-organized by:

- Providing enough embedded free space at database load or reorganization time. IMS can then use this free space to insert new segments near their related segments (segments in the same database record).

Related Reading: For details on how to provide enough embedded free space, see “Specifying Free Space (HDAM, PHDAM, HIDAM, and PHIDAM Only)” on page 245.

Tip: Choose the amount of free space based on the growth and performance characteristics of your database. For new databases, use a value of 25% and increase or decrease this value as needed. It is a good idea to schedule a reorganization for the database when the reusable free space is less than 5%.

- Selecting an appropriate database reorganization frequency.

Related Reading: For more information on when and how to reorganize your databases, see “Reorganizing the Database” on page 347.

- Using efficient HDAM and PHDAM randomizing modules and randomizing parameters. Information on this can be found in “Determining Which Randomizing Module to Use (HDAM and PHDAM Only)” on page 247.

Adjusting HDAM and PHDAM Options

To assess any design choices you have previously made or to improve performance, read “Choosing HDAM or PHDAM Options” on page 248. This topic discusses the HDAM and PHDAM options you can choose and the performance implications each.

You can adjust HDAM and PHDAM options using the reorganization utilities:

1. Determine whether the change you are making will affect the code in any application programs. It should only do so if you are changing to a sequential randomizing module.
2. Unload your database, using the existing DBD and the appropriate unload utility.
3. Code a new DBD (for non-PHDAM) using the TSO Partition Definition Utility. If you changed your CI or block size, you need to allocate buffers for the new size.

Related Reading: See Chapter 9, “Designing Full-Function Databases,” on page 245 for a discussion of what things to consider in choosing buffer number and size and how they are specified.

4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Determine whether you need to recalculate database space.

Related Reading: See “Estimating the Minimum Size of the Database” on page 315 for a description of how to calculate space.

6. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
7. Reload your database or partition using the new DBD (if any) and the appropriate reload utility. Make an image copy of your database as soon as it is reloaded.

Adjusting Buffers

The size and number of buffers you can choose are described in “Multiple Buffers in Virtual Storage” on page 253. This topic also discusses the performance implications of choosing a buffer size and number. To improve performance, reread that topic and reassess the original choices you made before you adjust your buffers.

VSAM Buffers

This topic contains the following information about VSAM buffers:

- “Monitoring VSAM Buffers” on page 412
- “When to Adjust VSAM Buffers” on page 412
- “VSAM Buffer Adjustment Options” on page 412

Monitoring VSAM Buffers

If you are using VSAM, you can monitor buffers using the DB monitor reports described in Chapter 14, “Monitoring Databases,” on page 341. For each buffer size you define, a VSAM subpool report is produced. The VSAM Buffer Pool report tells the number of buffers in the subpool and their size (in the SUBPOOL BUFFER SIZE and TOTAL BUFFERS IN SUBPOOL fields).

When to Adjust VSAM Buffers

Adjust VSAM buffers when you see buffer performance begin to degrade, or if you wish to add options to boost performance in anticipation of increased buffer activity.

VSAM Buffer Adjustment Options

1. If background write is turned on and the number in the NUMBER OF VSAM WRITES TO MAKE SPACE IN THE POOL field is not zero, you probably do not have enough buffers allocated in the subpool. Try allocating more buffers to decrease the number or reduce it to zero.
2. If you need to improve performance for a specific application, you can reserve subpools for certain data sets by:
 - Defining multiple local shared resource pools.
 - Dedicating subpools to a specific data set.
 - Defining separate subpools for index and data components of VSAM data sets. *IMS Version 9: Installation Volume 2: System Definition and Tailoring* tells you how to specify these options.
3. If sequential mode processing is not used, the number of VSAM buffers specified in the DFSVSAMP DD statement can dramatically affect performance. This problem occurs when the number of VSAM KSDS indexes that must be read, plus one for the data portion, is equal to or greater than the number of VSAM buffers allocated. This problem can be alleviated either by increasing the number of buffers or by using sequential mode. With sequential mode, the need to read indexes above the sequence set is reduced. However, sequential mode can only be obtained in a batch environment with a DBD referenced by a single PCB and with a processing option of LOAD or RETRIEVE only. Sequential mode is not available in data sharing.
4. VSAM buffers can take advantage of z/OS Hiperspace buffering.

Hiperspace Buffering Parameters: To use Hiperspace buffering, you must specify one or two optional parameters on the VSRBF subpool definition statement:

HSOIHSR

Specifies the action IMS takes if Hiperspace buffering requested for a subpool is unavailable.

HSO Hiperspace buffering is optional. IMS continues to run.

HSR Hiperspace buffering is required. IMS terminates.

HSn Specifies the number of Hiperspace buffers to build for a subpool. The number *n* is a 1- to 8-digit number.

Hiperspace parameters are valid only for buffer sizes of 4K or multiples of 4K. Specifying Hiperspace parameters on buffers smaller than 4K causes an error. To use Hiperspace buffering you might need to unload your database and then reload it into 4K or multiples of 4K CI sizes to accommodate Hiperspace requirements.

If you decide to leave intact databases with CI sizes of less than 4K, do not allocate any buffers less than 4K. The CIs that are less than 4K are placed in 4K or larger

buffer pools. However, the CIs compete with VSAM data sets already there. This method might be expedient in the short term.

Related Reading:

- For more information on coding the HSOIHSR and HSn parameters to activate Hiperspace buffering on VSAM buffers, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.
- For more information about VSAM buffers, including Hiperspace buffers, see *z/OS V1R4: DFSMS: Using Data Sets*.

OSAM Buffers

If you are using OSAM, individual subpool buffer reports do exist. However, you can monitor the number of buffers you are using by using the Enhanced OSAM Buffer Subpool statistics function which supports the following values:

DBESF

Provides the full OSAM Subpool statistics in a formatted form.

DBESU

Provides the full OSAM Subpool statistics in an unformatted form.

DBESS

Provides a summary of the OSAM database buffer pool statistics in a formatted form.

DBESO

Provides a the full OSAM database buffer pool statistics in a formatted form for online statistics returned as a result of a /DIS POOL command.

Related Reading: For detailed information on these values, see the *IMS Version 9: Application Programming: Design Guide*.

Another way to improve performance, this time for a specific application, is to reserve subpools for use by certain data sets. For example, if you have an index data set with a block size of 512 bytes, reserve a subpool for it that contains 512-byte buffers. You can do this by *not* defining 512-byte block sizes for any other data sets in the database. (Remember, block sizes are specified *by data set* in the BLOCK= operand in the DATASET statement in the DBD.) If you then allocate enough 512-byte buffers to hold all the blocks in your index, all blocks read into the buffer pool will remain in the buffer pool.

Performance can also be improved through the use of the **co** (caching option) parameter of the IOBF control statement specified either in the DFSVSMxxx member of IMS.PROCLIB or in DFSVSAMP.

Related Reading:

- For detailed information about the DB Monitor Database Buffer Pool report, see the *IMS Version 9: Utilities Reference: System*.
- For more information on the **co** (caching option) parameter of the IOBF control statement, OSAM buffer pools and the use of the coupling facility for OSAM data caching see the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Procedure for Adjusting VSAM and OSAM Database Buffers

To adjust VSAM and OSAM database buffers, change the control statements that specify buffer size and number. Then put the new control statements in the:

Adjusting Buffers

- DFSVSAMP data set in batch and utility environments
- IMS.PROCLIB data set with the member name DFSVSMnn in IMS TM and DBCTL environments

Related Reading: Detailed information on how to code these control statements is in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

OSAM Sequential Buffering

If you are using OSAM Sequential Buffering, you can use the Sequential Buffering Summary report and the Sequential Buffering Detail report to see how the SB buffers were used during a your program's execution.

By default, four buffer sets exist in each SB buffer pool. If the reports indicate that a large percentage of random read I/O operations were used, and you know that the program was processing your database sequentially, increasing the number of buffer sets to six or more can improve performance. By increasing the number of buffer sets, it is more likely that a block is still in an SB buffer when requested, and a read I/O operation is not necessary.

If only a few random reads were used during your program's execution, it indicates that the database is very well organized and most requests were satisfied from the SB buffer pool or with sequential reads. If this happens, you can save virtual storage space by decreasing the number of buffer sets in each SB buffer pool to two or three.

Procedure for Adjusting Sequential Buffers

You can change the number of buffer sets allocated to each SB buffer pool in two ways:

- Coding an SBPARM control statement with the BUFSETS keyword.
- Using an SB Initialization Exit Routine.

Once you have changed the number of buffer sets, you can use the SB Test Utility to reprocess the SB buffer handler call sequence that was issued during your program's execution. Then you can study the resulting //DFSSTAT reports to see the impact of the change.

Related Reading:

- The Sequential Buffering Summary report and the Sequential Buffering Detail reports are described and instructions on how to use the SB Test Utility are in the *IMS Version 9: Utilities Reference: Database and Transaction Manager*.
- Detailed instructions on how to code an SBPARM control statement are in the *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.
- Details on the SB Initialization Exit Routine are in the *IMS Version 9: Customization Guide*.

Adjusting VSAM Options

The VSAM options you can choose are described in "VSAM Options" on page 264. In Chapter 6, "Choosing Full-Function Database Types," on page 59, the performance implications of each VSAM option are also discussed. To improve performance, reread that topic and reassess the original choices you made.

The only VSAM option you can specifically monitor for is background write. If you are not using background write, you can look at the VSAM Buffer Pool report described in *IMS Version 9: Utilities Reference: System*. The report, in the Number of VSAM Writes To Make Space in the Pool field, documents the number of times data in a buffer had to be written to the database before the buffer could be used. If you use background write, you might find that you are able to reduce this number and therefore the size of the buffer pool.

If you are already using background write, the VSAM Buffer Pool report tells you how many times background write is invoked in the Number of Times Background Write Function Invoked field. The VSAM Statistics report (another report produced by the DB monitor) tells you in the BKG WTS field if background write was invoked. It also tells you, in the USR WRTS field, among other things, how many times background write was invoked.

Two types of adjustable VSAM options exist:

- Options specified in the OPTIONS control statement
- Options specified in the Access Method Services DEFINE CLUSTER command

Procedure for Adjusting VSAM Options Specified in the OPTIONS Control Statement

To adjust these VSAM options, change the appropriate parameters in the OPTIONS control statement. Then put the new control statement in the:

- DFSVSAMP data set in a batch system
- IMS.PROCLIB data set with the member name DFSVSMnn in an online system

Detailed information on how to code these control statements is in *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Procedures for Adjusting VSAM Options Specified in the Access Method Service DEFINE CLUSTER Command

To adjust these VSAM options, change the appropriate parameters in the DEFINE CLUSTER command. What additional things you must do depends on which VSAM parameter you are changing, as described in this topic.

Changing the FREESPACE Parameter

You can use the reorganization utilities to change the use of free space or to change the percent of free space you have specified. To make this change:

1. Unload your database using the existing DBD and the appropriate unload utility.
2. Recalculate database space. You need to do this because the change you are making will result in different requirements for database space. See Chapter 13, "Loading Databases," on page 315, "Estimating the Minimum Size of the Database" on page 315 for a description of how to calculate database space.
3. Delete the old database cluster and define the new database cluster with a change to the FREESPACE parameter.
4. Reload your database, using either the existing DBD (if no changes were made to the DBD) or the new DBD. Use the appropriate reload utility.
5. If the database being reorganized is a secondary index with direct pointers, you must run some of the reorganization utilities before and after reloading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Adjusting VSAM Options

Changing the SPEED / RECOVERY Parameter

Do not unload and reload your database merely to change the SPEED/RECOVERY parameter. Rather, if you have RECOVERY specified, change the parameter to SPEED to improve performance when the database is reloaded and restart of the load program is not used. IMS does not support the RECOVERY parameter. Recovery can only be done when the database load program is run under control of UCF.

Because it is assumed you would only change the parameter when making other database changes that require you to unload and reload your database, no procedure for changing it is provided here.

Changing the REPLICATE / NOREPLICATE Parameter

You can use the reorganization utilities to change whatever you've specified for the REPLICATE/NOREPLICATE parameters. To change them:

1. Unload your database, using the existing DBD and the appropriate unload utility.
2. Recalculate database space. You need to do this because the change you are making will result in different requirements for database space.

Related Reading: See Chapter 13, "Loading Databases," on page 315, "Estimating the Minimum Size of the Database" on page 315 for descriptions of how to calculate database space.

3. Delete the old database cluster and define the new database cluster.
4. Reload your database using the existing DBD and the appropriate reload utility.
5. If your database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after reloading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Adjusting OSAM Options

The OSAM options you can choose are described in "OSAM Options" on page 269. Performance implications of each OSAM option are also discussed there. To improve performance, reread that topic and reassess the original choices you made.

You cannot specifically monitor any OSAM options. To adjust OSAM options, change the appropriate parameters in the OPTIONS control statement. Then put the new control statement in the:

- DFSVSAMP data set in a batch system
- IMS.PROCLIB data set with the member name DFSVSMnn in an online system

Detailed information on how to code these control statements is in *IMS Version 9: Installation Volume 1: Installation Verification*.

Changing the Amount of Space Allocated

Change the amount of space allocated for your database in two situations. The first is when you are running out of primary space. Do not use your secondary space allocation because this can greatly decrease performance. Also change the amount of space allocated for your database when the number of I/O operations required to process a DL/I call is large enough to make performance unacceptable. Performance can be unacceptable if data in the database is spread across too much DASD space.

One way to routinely monitor use of space is by watching the IWAITS/CALL field in the DL/I Call Summary report. The DL/I Call Summary report is described in *IMS Version 9: Utilities Reference: System*. If the IWAITS/CALL field has a relatively high number in it, the high number can be caused by space problems. If you suspect space is the problem, you can verify such problems in two specific ways:

- For VSAM data sets, you can get a report from the VSAM catalog using the LISTCAT command. In the report, check CI/CA splits, EXCPs, and EXTENTS (LISTCAT ALL report is described in Chapter 14, “Monitoring Databases,” on page 341).
- For non-VSAM data sets, you can get a report on the VTOC using the LISTVTOC command. In the report, check the NOEXT field (LISTCAT ALL report is described in Chapter 14, “Monitoring Databases,” on page 341).

If you decide to change the amount of space allocated for your database, do it with JCL or with z/OS utilities. The reorganization utilities must be run to put the database in its new space. The procedure for putting the database in its new space is as follows:

1. Unload your database, using the existing DBD and the appropriate unload utility.
2. Recalculate database space.

Related Reading: See “Estimating the Minimum Size of the Database” on page 315 for a description of how to calculate database space.

3. Delete the old database space for non-VSAM data sets and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
4. If you are changing the space in the root addressable area of an HDAM database, you might need to adjust other HDAM parameters. In this case, you must code a new DBD before reloading (a new DBD is not needed when a PHDAM partition is changed). To change the space in the root addressable area of a PHDAM partition, you must use the HALDB Partition Definition utility.
5. Reload your database, using either the existing DBD (if no changes were made to the DBD) or the new DBD. Use the appropriate reload utility.
6. You must run some of the reorganization utilities before and after reloading to resolve prefix information if your non-HALDB database uses logical relationships or secondary indexes. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing Operating System Access Methods

You can use the reorganization utilities to change access methods from OSAM to VSAM, or from VSAM to OSAM.

To change access methods:

1. Unload the database.
2. Code a new DBD (unless you have already done this as described in Step 1).
3. Delete the old data sets and define the new clusters when changing from non-VSAM to VSAM. Delete the old clusters and define new database data sets when changing from VSAM to non-VSAM.
4. You need to change from OSAM options and buffers to VSAM options and buffers or vice versa. These topics are covered in preceding sections of this chapter:

“Adjusting Buffers” on page 411

“Adjusting VSAM Options” on page 414

“Adjusting OSAM Options” on page 416

5. Reload your database, using the new DBD. Remember to make an image copy of your database as soon as it is reloaded.
6. If your non-HALDB database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after loading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing the Number of Data Set Groups

Normally, a database is physically stored on one data set or, as in HISAM, on a pair of data sets. However, databases can be physically stored on more than one data set or pair of data sets. If so, each data set or pair of data sets is called a data set group. “Multiple Data Set Groups” on page 235 tells you:

- What data set groups are
- When they can be used
- What situations might prompt you to use them
- How they are specified in the DBD

You should be familiar with these topics. You should also have decided to change to multiple data set groups to tune your database. It is not possible for you to specifically monitor your database to determine whether multiple data set groups will improve performance or better utilize space. Rather, knowledge of your application’s requirements along with many types of statistics about database use might help you make this decision.

To change the number of data set groups in your database, (see Figure 225 on page 419) you:

1. Unload your database using the existing DBD.
2. If your database is PHDAM or PHIDAM, delete the database definition from the DBRC RECON data sets using the HALDB Partition Definition Utility.
3. Code a new DBD.
4. Recalculate database space. You need to recalculate database space because the change you are making will result in different requirements for database space.

Related Reading: See “Estimating the Minimum Size of the Database” on page 315 for a description of how to calculate database space.

5. Delete the old database space and define new database space for non-VSAM data sets. Delete the space allocated for the old clusters and define space for the new clusters for VSAM data sets.
6. If your new database is PHDAM or PHIDAM, run the HALDB Partition Definition utility to define the partition data sets for the database.
7. Reallocate data sets because the number and size of data sets you are using will change.

Related Reading: See “Allocating Data Sets” on page 322 for information on allocating data sets.

8. Reload your database using the new DBD. Take an image copy of your database as soon as the database is reloaded.
9. Run some of the reorganization utilities before and after reloading to resolve prefix information if your database uses logical relationships or secondary indexes. The flowchart in Figure 197 on page 352 shows you which utilities to use and the order in which they must be run.

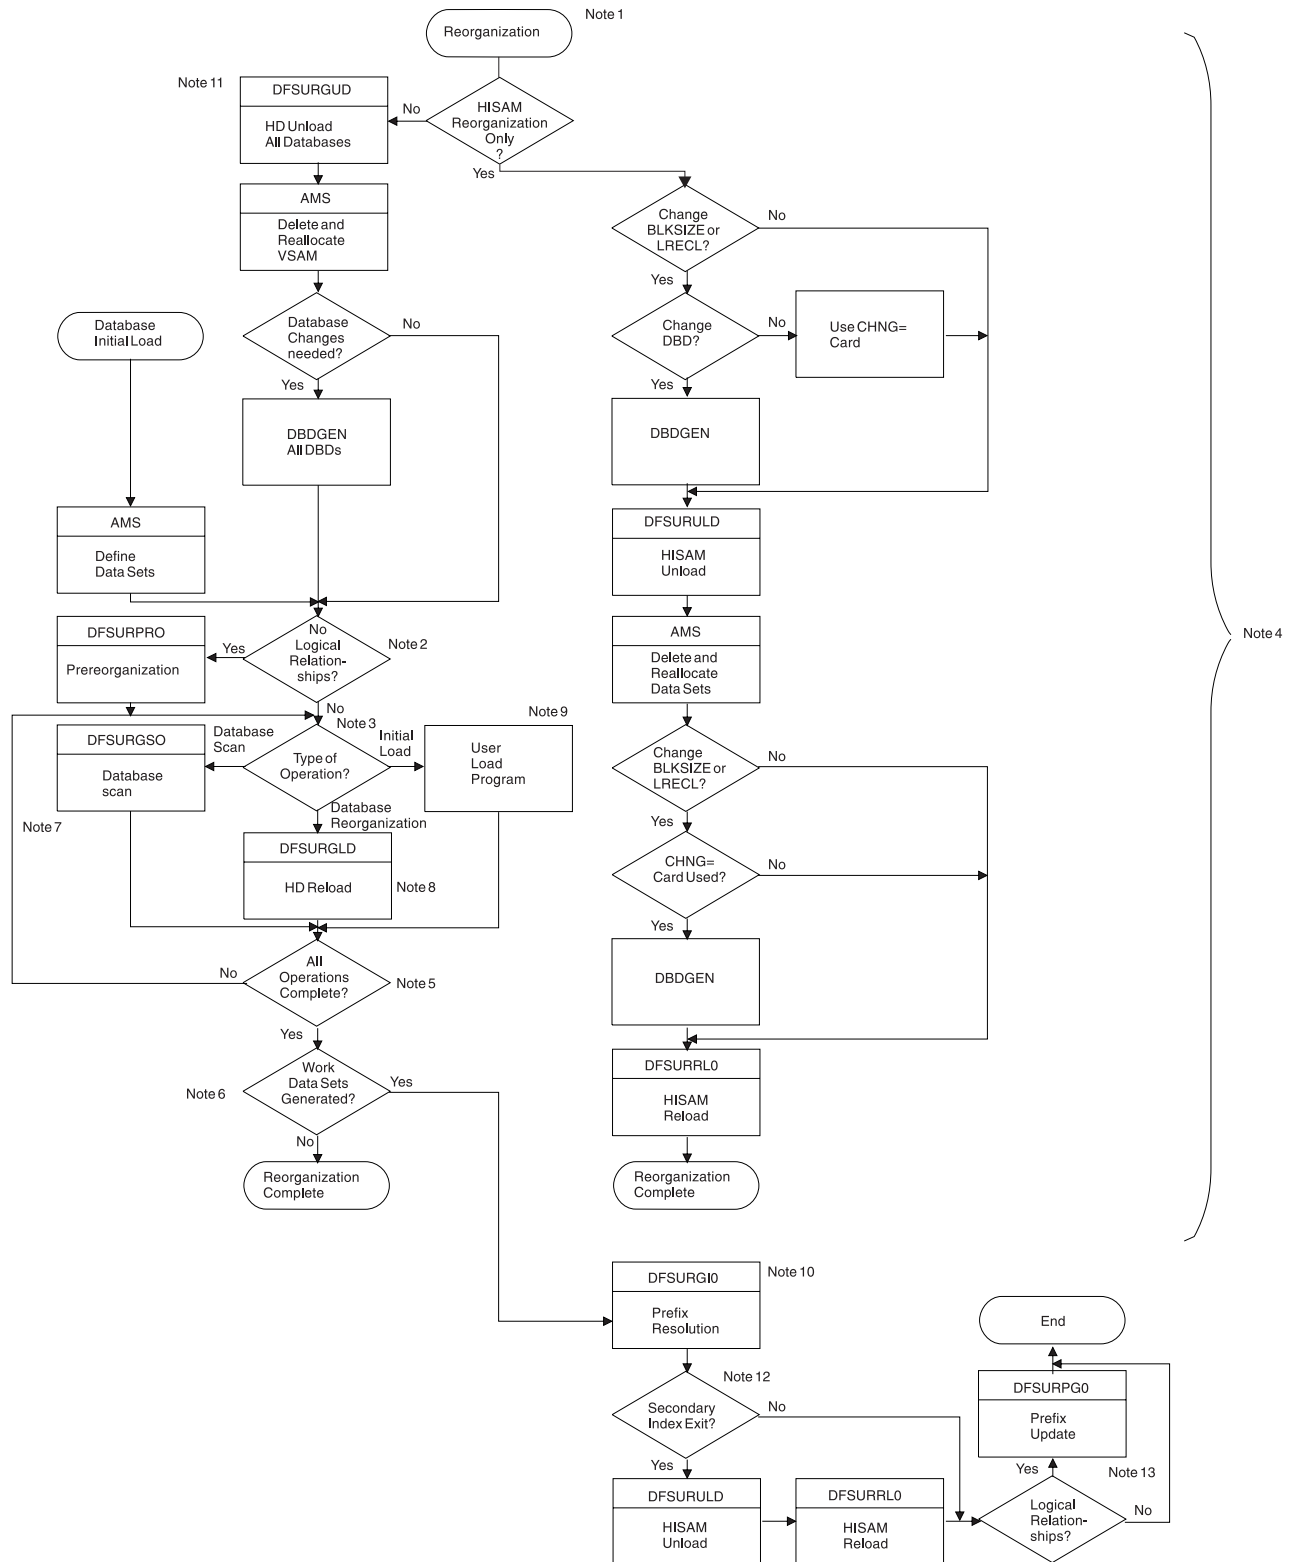


Figure 225. Utility Sequence of Execution When Making Database Changes during Reorganization

Notes to Figure 225:

1. You can use the database reorganization/load processing utilities (that is, the HISAM Unload/Reload, HD Unload/Reload, Prefix Resolution and Prefix

Changing the Number of Data Set Groups

Update utilities) to operate on one or more databases concurrently. For example, you can reorganize one or more existing databases at the same time that other databases are being initially loaded. Any or all of the databases being operated on can be logically interrelated. A database operation is defined as an initial database load, a database unload/reload (reorganization), or a database scan.

2. If one or more segments in any or all of the databases being operated upon is involved in either a logical relationship or a secondary index relationship, the YES branch must be taken. You can also use the Prereorganization utility to determine which database operations must be performed.
3. Based upon the information given to it on control statements, the database Prereorganization utility provides a list of databases that must be initially loaded, reorganized, or scanned. You must not change the number and sequence of databases specified on the prereorganization control statement between reload and prefix resolution.
4. This area of the flowchart must be followed once for each database to be operated upon, whether the operation consists of an initial load, reorganization, or scan. The operations can be done for all databases concurrently, or one database at a time. If the various database operations are performed sequentially, work data set storage space can be saved and processing efficiency increased if DISP=(MOD,KEEP) is specified for the DFSURWF1 DD statement associated with each database operation. The attributes of the work data set for the database initial load, reorganization, and scan programs must be identical.

When using the HD Reorganization Reload utility, first do all unloads and scans of logically related databases if logical parent concatenated keys are defined as virtual in the logical child.

5. You must ensure that all operations indicated by the Prereorganization utility (if it was executed) are completed prior to taking the YES branch.
6. If any work data sets were generated during any of the database operations that were executed by you, the YES branch must be taken. The presence of a logical relationship in a database does not guarantee that work data sets will be generated during a database operation. The reorganization/load processing utilities determine the need for work data sets dynamically, based upon the actual segments presented during a database operation. If any segments that participate in a logical relationship are loaded, work data sets will be generated and the YES branch must be taken.

If for any specific database operation no work data set was generated for the database, processing of that database is complete and ready to use.

When a HIDAM database is initially loaded or reorganized, its primary index will be generated at database load time.

7. You must run the DB Scan utility before a database is unloaded when logical parent concatenated keys are defined as virtual in the logical child database to be unloaded.

This program should be executed against each database listed in the output of the Prereorganization utility. A work data set can be generated for each database scanned by this utility. Databases for scanning are listed after the characters "DBS=" in one or more output messages of the Prereorganization utility.

8. The HD Reorganization Reload utility can cause the generation of a work data set to be later used by the Prefix Resolution utility. Databases to be reorganized using the HD Reorganization Unload utility and the HD

Reorganization Reload utility are listed after the character “DBR=” in one or more output messages of the Prereorganization utility.

9. The user-provided initial database load program can automatically cause the generation of a work data set to be later used by the Prefix Resolution utility. You do not need to add code to the initial load program for work data set generation. Code is added automatically by IMS through the user program issuing ISRT requests. You must, however, provide a DD statement for this data set along with the other JCL statements necessary to execute the initial load program. Databases for initial loading are listed after the characters DBIL= in one or more output messages of the Prereorganization utility.
10. The database Prefix Resolution utility combines the workfile output from the Database Scan utility, the HD Reorganization Reload utility, and the user's initial database load execution to create an output data set for use by the Prefix Update utility. The Prefix Update utility then completes all logical relationships defined for the databases that were operated upon.
11. This path must be taken for HISAM databases with logical relationships. This path must also be taken if structural changes are required (for example, HISAM to HDAM, pointer changes, additional segments, or adding a secondary index).
12. If a secondary index needs to be created or if two secondary indexes need to be combined, you must run the HISAM Unload/Reload utilities. After the HISAM Unload/Reload utilities are run, if logical relationships exist in the database, you must execute the Prefix Update utility before the reorganization or load process is considered to be complete.
13. For information on scratching and allocating OSAM data sets, see the topic about designing the IMS online system in *IMS Version 9: Administration Guide: System*.

Tuning Fast Path Systems

Your objective in tuning the IMS online system when Fast Path applications are present depends upon the importance of the message-driven programs and their criteria for acceptable response time. The performance analysis studies that you should undertake are:

- Examining the availability of sufficient real storage
- Checking the effectiveness of the balancing groups
- Investigating the number of Fast Path dependent regions and the possibility of parallel processing
- Monitoring of the required frequency of DEDB reorganization to reduce fragmented units of work
- Monitoring of the use of DEDB overflow buffers
- Monitoring the forced serialization of programs that concurrently need to use overflow buffers specified by the EXEC statement DBFX parameter
- Examining the area key ranges and whether the randomizing algorithm can be refined
- Reducing the amount of mixed mode processing

Fast Path performance can also be improved by eliminating unnecessary delays caused by the following:

- Transaction volume to a particular Fast Path application program
- DEDB structure considerations
- Contention for DEDB Control Interval (CI) resources

Tuning Fast Path Systems

- Exhaustion of DEDB DASD space
- Utilization of available real storage
- Sync point processing and physical logging
- Contention for output threads (OTHR)
- Overhead resulting from reprocessing
- Dispatching priority of processor-dominant and I/O-dominant tasks
- DASD contention caused by I/O on DEDBs
- Resource locking considerations with block level sharing
- Buffer pool usage and not grouping Fast Path application programs with similar buffer use characteristics together into one or more message classes

Statistics on transaction processing and contention for CIs can be obtained from the output of the Fast Path Log Analysis utility (DBFULTA0), which retrieves (from system log input) data relating to the usage of Fast Path resources.

Related Reading: For information on the Fast Path Log Analysis utility, see *IMS Version 9: Utilities Reference: System*.

Transaction Volume to a Particular Fast Path Application Program

If a disproportionately high number of transactions are queued to a particular balancing group, consider increasing the number of regions associated with that particular balancing group. The Fast Path Log Analysis report provides information about balancing group queuing.

DEDB Structure Considerations

Several characteristics of DEDB usage affect an application's response time:

- Data replication
- Subset pointers
- Number of areas
- Complexity of hierarchic structure
- Complexity of DL/I calls
- Use of sharing across IMS
- Last child pointers
- Recoverability

The first three characteristics are unique to DEDBs; the last five apply generally to databases. Data replication allows up to seven data sets for an individual area. When reading from an area represented by multiple data sets, performance is not impacted, unless the CI is defective. When updating, up to seven additional writes could be required. Although the physical write is performed asynchronously to transaction processing, there could be delays caused by access paths to a variety of DASD devices.

Up to eight subset pointers allow an application program to separate the children of a parent into groups in a DEDB, with the subset pointer pointing to the start of each group. Use of such pointers can help improve performance by reducing the time needed to access segments whose position is significantly displaced in a chain of sequential dependent segments.

Usage of Buffers from a Buffer Pool

The Fast Path buffer pool is used by all Fast Path programs except the DEDB online utilities, which have their own buffer pool. The Fast Path buffer pool is used to support the processing of MSDBs and DEDBs. The Fast Path buffer pool comprises buffers of a size defined at system startup by the BSIZ parameter. The buffer size selected must be capable of holding the largest CI from any DEDB area that is to be opened. The number of buffers page-fixed is based upon the value of supplied parameters:

- The normal buffer allocation (NBA) value causes the defined number of buffers to be fixed in the buffer pool at startup of the dependent region. (This number can be specified for the dependent region startup procedure using the NBA parameter.) The application program in this dependent region is eligible to receive up to this number of buffers within a given sync interval before one of the following occurs:
 - The buffer manager acquires unmodified buffers from the requesting application program.
 - No more buffers can be acquired on behalf of the requesting application program (a number of buffers equal to NBA have been requested, received, and modified). In this case, the buffer manager must acquire access to the overflow buffer allocation (OBA) if this value was specified for this program. If no OBA was specified, then all resources acquired for this program during sync interval processing to date are released.
- The OBA value is the number of buffers that a program can serially acquire when NBA is exceeded. (This number can be specified for the dependent region startup procedure using the OBA parameter.) The overflow interlock function serializes the overflow buffer access, and only one application program at a time can gain access to the overflow buffer allocation. Therefore, the overflow buffer can be involved in deadlocks.
- The DBFX value, which is a system startup parameter, defines a reserve of buffers that are page-fixed upon start of the first Fast Path application program. These buffers are used when asynchronous OTHREAD processing is not releasing buffers quickly enough to support the requests made in sync interval processing.

It follows that:

- BSIZ should be set equal to the largest DEDB CI that will be online. Because the buffer manager does not split buffers to accommodate multiple control intervals, making all DEDB CIs of a same size will provide more optimum use of storage. Even though large block sizes (up to 28K) can be used, this would cause only partial use of the buffer pool if there were many smaller CI sizes.
- The NBA value should be set approximately equal to the normal number of buffer updates made during a sync interval. The NBA value for inquiry-only programs should be small, because the buffers that are never modified can be reused and will all be released at sync time.
- The OBA should be used only in relation to a limited proportion of sync intervals. OBA is not required for inquiry-only programs. In general, the user should be careful to use the OBA value as intended. It should be used to support sync intervals where application program logic demands a variation in total modified buffer needs, thereby requiring access to OBA on an exceptional basis. With BMPs, OBA values greater than 1 should be unnecessary because the 'FW' status code that is returned when the NBA allocation is exceeded can be used to invoke a SYNC call. Invoking a SYNC call would then release all resources.

Tuning Fast Path Systems

Such application design reduces the serialization and possible deadlocks inherent in using the overflow interlock function.

- The DBFX value should be set, taking into account the total number of buffers that are likely to be in OTHREAD processing at peak load time. If this value is too low, an excessive number of wait-for-buffer conditions are reflected in the IMS Fast Path Log Analysis report.

To optimize the buffer usage, group message processing application programs with similar buffer use characteristics and assign them to a particular message class, so that the applications share the region's buffers.

Related Reading: See *IMS Version 9: Installation Volume 2: System Definition and Tailoring* for details of APPLCTN and TRANSACT class specifications.

Contention for DEDB Control Interval (CI) Resources

Queuing takes place on the DEDB CI resource to maintain serialized access on DEDB data. When two independent application programs concurrently request access to a particular CI, one requestor is required to wait. When such a wait would cause a deadlock, one of the application programs is selected to have its resources released and its processing returned to the previous sync point. (It should be noted that the overflow buffer interlock can also be involved in a deadlock). The rules for selection of the program to be interrupted because of a deadlock are:

- If the deadlock involves one or more message-driven programs, one of the programs is abnormally terminated, reinstated to its previous sync point, and rescheduled.
- If a BMP deadlocks with another BMP, the BMP that went through sync point last is abnormally terminated, has its resources released, is sent back to its previous sync point, and is given a return code.
- If a deadlock involves a DEDB utility, the other program is terminated and rescheduled. Two utilities cannot be involved in a deadlock, because two utilities cannot concurrently access the same DEDB area.

The number of contention and deadlock situations can be decreased by taking the following steps:

- Ensure that CIs contain no more segments than necessary. (CI size is specified in the DBD.)
- Limit the use of the overflow buffer interlock, which, in conjunction with CI usage, can be involved in a deadlock.
- Limit the value of NBA to the value necessary to cope with the majority of cases and use OBA to deal with the exceptional conditions. When the full buffer allocation (NBA or NBA and OBA) for a program has been exceeded, the buffer manager can begin stealing unmodified buffers from this program. When all buffers associated with a CI have been stolen, the CI can be released, providing it is not currently in use by a PCB. The buffer stealing and associated CI releasing is triggered by exceeding the full buffer allocation. Minimizing NBA and OBA will assist the timely release of CIs, thereby reducing CI contention.
- Ensure that BMPs accessing DEDBs issue SYNC calls at frequent intervals. (BMPs could be designed to issue many calls between sync points and so gain exclusive control over a significant number of CIs.)
- BMPs that do physical-sequential processing through a DEDB should issue a SYNC call when crossing a CI boundary (provided it is possible to calculate this point). This ensures that the application program never holds more than a single CI.

Reports produced by the Fast Path Log Analysis utility give statistics about CI contention.

Exhaustion of DEDB DASD Space

An out-of-space condition (with consequent stoppage of the DEDB area) can occur in the root addressable and sequential dependent portions of an area. Such situations will affect the operation of the system as a whole and can necessitate lengthy recovery procedures. The number of out-of-space conditions can be decreased by:

- Attempting to restrict the number of uses of independent overflow CIs through randomizing algorithm design or regular reorganization
- Deleting sequential dependent CIs on a regular basis
- Using display commands or DEDB POS calls to track space usage

An out-of-space condition can be relieved without bringing IMS down by following the procedures in “Extending DEDB Independent Overflow Online” on page 464.

Utilization of Available Real Storage

The amount of page-fixed storage defined will be a significant consideration in limited storage systems. The factors influencing real storage utilization are summarized in Appendix B, “Insert, Delete, and Replace Rules for Logical Relationships,” on page 471.

Synchronization Point Processing and Physical Logging

Some 'clustering' of output and release of updated CIs and buffers occurs because DEDB updates are deferred until after physical logging is complete. In BMPs, it helps to minimize the number of updates performed in any one sync interval, particularly if the program is to be run concurrent with the main bulk of message processing.

It is likely that, for performance reasons, the physical log record will be large, so that the log record might not be written for some time during low logging activity. However, IMS varies the interval between the periodic invoking of physical logging. This interval is directly related to the total logging activity in the IMS system. (Low activity causes a smaller interval to be set.)

The physical logging process can be relatively slow because of small physical log buffers or channel or control unit contention for the WADS/OLDS data sets.

The Fast Path environment can have high transaction rates and logging activity. Therefore, the physical configuration supporting the logging process must also be analyzed and altered for optimum performance.

Contention for Output Threads

Each OTHR defined provides for the possibility of scheduling a separate service request block (SRB) to control the writing of the modified buffers associated with a particular sync interval. If the OTHR value is low, then queuing of write buffers waiting for an output thread can occur. In general, it is probably best to have one OTHR for each started dependent region that will cause modification of a DEDB.

Overhead Resulting from Reprocessing

Overhead will result from the necessity to perform reprocessing in either the message-driven or non-message-driven environments. The following conditions will necessitate reprocessing:

- Deadlocks involving CIs and (possibly) overflow interlock
- Verify failures at sync point time
- User-initiated rollback caused by such conditions as verify failure at call time

In the case of deadlocks, the application program is pseudo abended for dynamic backout. The program controller subtask is detached, and subsequently, reattached. For verify failures or rollback calls, rescheduling involves only the release of resources held and returned to the application program.

Excessive incidence of the above conditions will add to response time and total overhead. Conditions resulting in abend interception followed by dump and application program reinstatement will add to overhead.

Dispatching Priority of Processor-Dominant and I/O-Dominant Tasks

Because MSDB processing within a sync interval is processor-dominant, application programs processing solely or mainly MSDBs should be dispatched at a lower priority than those programs processing solely or mainly DEDBs (I/O dominant).

DASD Contention Due to I/O on DEDBs

As always, I/O contention for DEDB Areas will act as a limitation upon performance. To minimize this impact:

- Limit the number of heavily-used Areas per device.
- Limit the number of application programs accessing any one DEDB area. One possibility here is to design the transaction, input edit/routing exit, and randomizing algorithm combination so that the access to any one area is limited to a particular application program or programs.
- Limit the incidence and effect of stealing unmodified buffers by appropriate application program design. Buffer stealing can necessitate a second I/O to recover the stolen buffer/control interval. This can happen if the logic of the application program requires processing of a buffer when a significant number of calls have been made following the first retrieval.

Resource Locking Considerations with Block Level Sharing

Resource locking can occur either locally in a non-sysplex environment or globally in a sysplex environment.

In a non-sysplex environment, local locks can be granted in one of three ways:

- **Immediately** because of either of the following reasons:
 - IMS was able to get the required IRLM latches, and there is no other interest on this resource.
 - The request is compatible with other holders or waiters.
- **Asynchronously** because the request could not get the required IRLM latches and was suspended. (This can also occur in a sysplex environment.) The lock is granted when latches become available and one of two conditions exist:
 - No other holders exist.
 - The request is compatible with other holders or waiters.

- **Asynchronously** because the request is not compatible with the holders or waiters and was granted after their interest was released. (This could also occur in a sysplex environment.)

In a sysplex environment, global locks can be granted in one of three ways:

- **Locally by the IRLM** because either of the following two reasons:

There is no other interest for this resource.

This IRLM has the only interest, this request is compatible with the holders or waiters on this system, and XES already knows about the resource.

- **Synchronously on the XES CALL** because:

Either XES shows no other interest for this resource.

Or XES shows only SHARE interest for the hash class.

- **Asynchronously on the XES CALL** because of one of two conditions:

Either XES shows EXCLUSIVE interest on the hash class by an IRLM, but the resource names do not match (FALSE CONTENTION by RMF).

Or the request is incompatible with the other HOLDERS and is granted by the CONTENTION Exit after their interest is released (IRLM REAL CONTENTION).

Resource Name Hash Routine

The Fast Path Resource Name Hash routine generates the hash value used by the IRLM. You can specify the name of such a routine with the USRHASH parameter on the FPCTRL macro, but it is ignored.

One technique used by the IMS-supplied Fast Path Resource Name Hash routine (DBFLHSH0) increases the range of values implicit with the relative CI numbers by combining parts of the 31-bit CI number with values derived from a database's DMCB number and its area number as follows: Bits 11 through 15 of DMCB number are XOR'd with bits 7, 6, 5, 4, 3 of the area number to give a combination 5-bit position number. (Using the area number's bits in reverse order helps make both DMCB number and area number vary the combination value.)

For the relative CI number (bits 0 through 15 are not used):

- Bits 16 through 20 are XOR'd with the combination value.
- Bits 21 through 25 are XOR'd with the combination value.
- Bits 26 through 29 are used unchanged.
- Bits 30 and 31 are not used—thus a hashed CI number used as a GHT entry represents four CIs.

For the hashed resource name:

- Bits 16 through 29 of the hashed relative CI become bits 18 through 31 of the hash value that is passed to the IRLM.
- Bits 18 through 26 of the hash value are used as the displacement into the resource hash table (RHT).
- Bits 18 through 31 are used as the displacement into the GHT.

Chapter 16. Modifying Databases

Under several circumstances, you must modify your database. Over time, user requirements can change, necessitating changes in the database design. Or you might choose to use new or different options or features. Or perhaps you have simply found a more efficient way to structure the database. This chapter describes the various types of structural changes you can make to your database and tells you when and how you can make the changes using the reorganization utilities.

This chapter examines the following areas of modifying a database:

- “Adding Segment Types” on page 430
- “Deleting Segment Types” on page 431
- “Moving Segment Types” on page 432
- “Changing Segment Size” on page 432
- “Changing Data in a Segment (Except for Data at the End of a Segment)” on page 433
- “Changing the Position of Data in a Segment” on page 433
- “Adding Logical Relationships” on page 433
- “Adding a Secondary Index” on page 451
- “Adding or Converting to Variable-Length Segments” on page 451
- “Converting to the Segment Edit/Compression Exit Routine” on page 452
- “Converting Databases for Data Capture Exit Routines and Asynchronous Data Capture” on page 453
- “Converting a Logical Parent Concatenated Key from Virtual to Physical or Physical to Virtual” on page 454
- “Using the Online Change Function” on page 454
- “Extending DEDB Independent Overflow Online” on page 464

When you modify your database, you often make more than a simple change to it. For example, you might need to add a segment type and a secondary index. This topic has procedures to guide you through making each type of change. If you make more than one change at a time, you should look at Figure 225 on page 419. The flowchart, when used with the individual procedures in this chapter, will guide you in making some types of multiple changes to the database.

Attention: If the DBD for an existing MSDB is changed, the header information (BHDR) might change, even though the database segments do not. In this case, the headers in the MSDBCPx data sets are invalid or the wrong length. A change in the MSDB headers causes message DFS2593I. If ABND=Y is specified in the MSDB PROCLIB member, ABENDU1012 is also issued. Correct this problem by using the MSDBLOAD option on a warm start or cold start to load the MSDBs from an MSDBINIT data set.

Related Reading: If you share data, additional information about modifications is in *IMS Version 9: Administration Guide: System*.

Adding Segment Types

There are three ways to add a segment type to a database:

- Unloading and reloading using the reorganization utilities
- Without unloading or reloading
- Using your own unload and reload program

Unloading and Reloading Using the Reorganization Utilities

You can add segment types to a database record using the reorganization utilities if:

- The segment type to be added is at the bottom level of a path in the hierarchy. Figure 226 shows an existing database record (indicated by solid lines) and the places where a new segment type can be added (indicated by dashed lines).
- The existing relative order of segments in the database record does not change. In other words, the existing parent to child relationships cannot change.
- The existing segment names do not change.

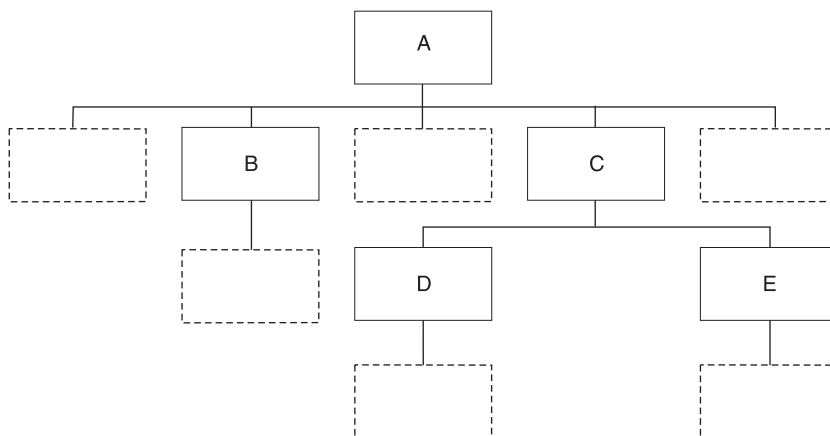


Figure 226. Where Segment Types Can Be Added in a Database Record

To use the reorganization utilities to add a segment type to the database:

1. Determine if the change you are making affects the code in any application programs. If the code is affected, make the necessary changes to the application program.
2. Unload your database, using the existing DBD.
3. Code a new DBD. You need to add SEGM= statements to the DBD for the new segment type. No database updates are allowed between unload and reload.
4. If the change you are making affects the code in application programs, make any necessary changes to the PSBs for those application programs.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space. You need to do this because the change you are making will result in different requirements for database space.

Related Reading: See “Estimating the Minimum Size of the Database” on page 315 for a description of how to calculate database space.

7. For non-VSAM data sets, delete the old database space and define the new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Reload your database, using the new DBD. Make an image copy of your database as soon as it is reloaded.
9. If your database uses logical relationships or secondary indexes, run some of the reorganization utilities before and after reloading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.
10. Code and execute an application program to insert the new segment types into the database.

Without Unloading or Reloading

You can add segment types to a database record without unloading the database under the following circumstances:

- In a HISAM database, the segment type to be added must be the last segment in the hierarchy. In addition, the segment type to be added must fit in the existing logical record.
- In an HD database, the segment type to be added must also be the last segment in the hierarchy. The parent of the new segment type must use hierarchic pointers. Also, the segment type cannot be the largest segment type in the data set group.

To add a segment type to the database without unloading and reloading:

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Code a new DBD. You need to add a SEGM= statement to the DBD for the new segment type.
3. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
4. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
5. Code and execute an application program to insert the new segment type.

Using Your Own Unload and Reload Program

You must write your own unload and reload program to add a segment type to the database, if the segment type does not meet the qualifications described in “Unloading and Reloading Using the Reorganization Utilities” on page 430 and “Without Unloading or Reloading.”

Deleting Segment Types

You can delete a segment type from a database by:

- Using the reorganization utilities
- Using your own unload and reload program

You can delete a segment type from a database, using the reorganization utilities, if:

- The existing relative order of segments in the database record does not change. In other words, the existing parent to child relationships cannot change.
- The existing segment names do not change.

Deleting Segment Types

To use the reorganization utilities to delete a segment type from the database:

1. Code and execute an application program to delete all occurrences of the segment type being deleted. You must code and execute the application program before the database is unloaded.
2. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
3. Unload your database, using the existing DBD.
4. Code a new DBD. You need to remove SEGM= statements from the DBD for:
 - The segment type being deleted
 - The children of the deleted segment.
5. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
6. Recalculate database space. You need to do this because the change you are making will result in different requirements for database space.
Related Reading: See “Estimating the Minimum Size of the Database” on page 315 for a description of how to calculate database space.
7. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
8. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
9. Reload your database using the new DBD. Remember to make an image copy of your database as soon as it is reloaded.
10. If your database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after reloading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Moving Segment Types

Because segment types cannot be moved using the reorganization utilities, you must write your own unload and reload program to move them.

Changing Segment Size

Using the reorganization utilities, you can increase or decrease segment size at the end of a segment type. When increasing segment size, you are adding data to the end of a segment. When decreasing segment size, IMS truncates data at the end of a segment.

If you are increasing the size of a segment, you cannot predict what is at the end of the segment when it is reloaded. Also, new data must be added to the end of a segment using your own program after the database is reloaded.

To increase or decrease segment size:

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload your database, using the existing DBD. If you are changing a HISAM database, you must use the HD UNLOAD/RELOAD utility since the HISAM utilities cannot be used to make structural changes.

3. Code a new DBD. You need to change the BYTES= operand on the SEGM statement in the DBD to reflect the new segment size. If you are eliminating data from a segment for which FIELD statements are coded in the DBD, you need to eliminate the FIELD statements. If you are adding data to a segment and the data is referenced in the SSA in application programs, you need to code FIELD statements. No database updates are allowed between unload and reload.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than build dynamically.
6. Recalculate database space. You need to do this because the change you are making results in different requirements for database space.

Related Reading: See “Estimating the Minimum Size of the Database” on page 315 for a description of how to calculate database space.

7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Reload your database, using the new DBD. Make an image copy of your database as soon as it is reloaded.
9. If your database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after reloading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Changing Data in a Segment (Except for Data at the End of a Segment)

Data in a segment cannot be increased or decreased in size using the reorganization utilities. To increase or decrease the size of fields, you must write your own unload and reload programs.

Changing the Position of Data in a Segment

You cannot change the position of data in a segment using the reorganization utilities. To make this kind of change, you must write your own unload and reload program, use field-level sensitivity, or use a tool, such as the IMS High Performance Pointer Checker for z/OS, which includes a DB Segment Restructure utility.

Related Reading:

- For information about how to use field-level sensitivity, see “Field-Level Sensitivity” on page 226
- For information about the IMS High Performance Pointer Checker for z/OS, see *IBM IMS High Performance Pointer Checker for z/OS User's Guide Volume 2*

Adding Logical Relationships

Logical relationships are explained in detail in “Logical Relationships” on page 157. This topic contains examples and procedures for adding a logically-related database to an existing database. Not all situations in which you might need to add a logical relationship are described in this topic. However, if the examples do not fit your specific requirements, you should be able to gather enough information from them to decide:

Adding Logical Relationships

- If adding a logical relationship to your existing database is possible
- How to add the relationship

The examples in this topic are followed by Table 32 on page 447, which tells you what to do when reorganizing a database to add a logical relationship. Following the table, “Some Restrictions on Modifying Existing Logical Relationships” on page 449 discusses some restrictions on modifying existing logical relationships.

The examples in this topic show the logical parent as a root segment, although this is not a requirement. The examples are still valid when the logical parent is at a lower level in the hierarchy.

When adding logical relationships to existing databases, you should always make the change on a test database. Thoroughly test the change before implementing it using production databases.

In the following examples, these conventions are used:

- Existing databases are shown using solid lines.
- The database being added is shown using dashed lines.
- The logical parent and logical child relationship is labeled for the database being added. They are labeled LP and LC.
- The terms DBX, DBY, and DBZ refer to database 1, database 2, and database 3.

Related Reading: For example procedures 1 through 13, the following related readings provide more detailed information for some of the steps:

- See “Estimating the Minimum Size of the Database” on page 315 for a description of how to calculate database space.
- See “Writing a Load Program” on page 325 for a description of how to write an initial load program.

Example 1. DBX Exists, DBY Is to Be Added

Example 1 is shown in Figure 227.




Figure 227. DBX Exists, DBY Is to Be Added

DBX must be reorganized to add the counter field to the segment prefix for A. DBIL must be specified in the control statement for DBX. In the following “Example 1 Procedure,” the counter field for segment A is updated to show the number of C segments because segment C is loaded with a user load program.

Example 1 Procedure

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.

2. Unload DBX, using the existing DBD and the HD Unload utility.
3. Code a new DBD for DBX and DBY. “How to Specify Use of Logical Relationships in the Logical DBD” in Chapter 8, “Choosing Optional Database Functions,” on page 157, explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBX and calculate space for DBY.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Prereorganization utility, specifying DBIL in the control statements for DBX and DBY.
9. Reload DBX, using the new DBD and the HD Reload utility.
10. Load DBY, using an initial load program.
11. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9 and 10 as input.
12. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 11 as input.
13. Remember to make an image copy of both databases as soon as they are loaded.

Example 2. DBX and DBY Exist, DBZ Is to Be Added

Example 2 is shown in Figure 228.




Figure 228. DBX and DBY Exist, DBZ Is to Be Added

In this example, the counter exists in the segment C prefix. DBX and DBY must be reorganized to calculate the new value for the counter in the segment C prefix. DBIL must be specified in the control statement for DBX and DBY. In the following “Example 2 Procedure,” the segment A counter field is updated to show the number of C segments because segment C is loaded with a user load program.

Example 2 Procedure

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBX and DBY, using the existing DBDs and HD Unload utility.

Adding Logical Relationships

3. Code a new DBD for DBY and DBZ. “How to Specify Use of Logical Relationships in the Logical DBD” in Chapter 8, “Choosing Optional Database Functions,” on page 157 explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBX and DBY, and calculate space for DBZ.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Preorganization utility, specifying DBIL in the control statements for DBX, DBY and DBZ.
9. Reload DBX and DBY, using the new DBDs and the HD Reload utility.
10. Load DBZ, using an initial load program.
11. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9 and 10 as input.
12. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 11 as input.
13. Remember to make an image copy of all three databases as soon as they are loaded.

Example 3. DBX and DBY Exist, DBZ Is to Be Added

Example 3 is shown in Figure 229.




Figure 229. DBX and DBY Exist, DBZ Is to Be Added

DBY must be reorganized to add the counter field to the segment C prefix. DBIL must be specified in the control statement for DBY. DBX must be reorganized because an initial load (DBIL) of the logical parent (segment C) assumes an initial load (DBIL) of the logical child. The procedure for this example (and all conditions and considerations) is exactly the same as example 2.

Example 4. DBX and DBY Exist, DBZ Is to Be Added

Example 4 is shown in Figure 230 on page 437.




Figure 230. DBX and DBY Exist, DBZ Is to Be Added

The procedure for this example (and all conditions and considerations) is exactly the same as for example 2.

Example 5. DBX Exists, DBY Is to Be Added

Example 5 is shown in Figure 231.




Figure 231. DBX Exists and DBY Is to Be Added

DBX must be reorganized to add the logical child pointers in the segment A prefix.

Procedure

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBX, using the existing DBD and the HD Unload utility.
3. Code a new DBD for DBX and DBY. "How to Specify Use of Logical Relationships in the Logical DBD" in Chapter 8, "Choosing Optional Database Functions," on page 157 explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBX, and calculate space for DBY.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Prereorganization utility, specifying DBR in the control statement for DBX, and DBIL in the control statement for DBY.
9. Reload DBX, using the new DBD and the HD Reload utility.
10. Load DBY, using an initial load program.

Adding Logical Relationships

11. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9 and 10 as input.
12. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 11 as input.
13. Remember to make an image copy of both databases as soon as they are loaded.

Example 6. DBX and DBY Exist, DBZ Is to Be Added

Example 6 is shown in Figure 232.




Figure 232. DBX and DBY Exist, DBZ Is to Be Added

DBY must be reorganized to add the logical child pointers to the segment C prefix. One of the following three procedures should be used:

- “Procedure When Reorganizing DBY (Segment B Contains a Symbolic Pointer)”
- “Procedure When Reorganizing DBY and Scanning DBX (Segment B Contains a Direct Pointer)” on page 439
- “Procedure When Reorganizing DBX and DBY” on page 439

Procedure When Reorganizing DBY (Segment B Contains a Symbolic Pointer)

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBY, using the existing DBD and HD Unload utility.
3. Code a new DBD for DBY and DBZ. “How to Specify Use of Logical Relationships in the Logical DBD” in Chapter 8, “Choosing Optional Database Functions,” on page 157 explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBY, and calculate space for DBZ.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Preorganization utility, specifying DBR in the control statement for DBY, and DBIL in the control statement for DBZ. (The output from the Preorganization utility indicates that a scan of DBX is required.)
9. Reload DBY, using the new DBD and the HD Reload utility.
10. Load DBZ, using an initial load program.

11. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9 and 10 as input.
12. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 11 as input.
13. Remember to make an image copy of both databases as soon as they are loaded.

When DBY is reloaded, two type 00 records are produced for each occurrence of segment C. One contains a logical child database named DBZ and matches the type 10 record produced for segment E. The other contains a logical child database named DBX. Because a scan or reorganization of DBX was not done, a matching 10 record was not produced for segment B. The Prefix Resolution utility produces message DFS878 when this occurs. The message can be ignored as long as the printed 00 record refers to DBY and DBX. Any messages for DBY and DBZ should be investigated.

Procedure When Reorganizing DBY and Scanning DBX (Segment B Contains a Direct Pointer)

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBY, using the existing DBD and HD Unload utility.
3. Code a new DBD for DBY and DBZ. "How to Specify Use of Logical Relationships in the Logical DBD" in Chapter 8, "Choosing Optional Database Functions," on page 157 explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBY, and calculate space for DBZ.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Prereorganization utility, specifying DBR in the control statement for DBY, and DBIL in the control statement for DBZ. (The output from the Prereorganization utility says that a scan of DBX is required.)
9. Run the scan utility against DBX.
10. Reload DBY, using the new DBD and the HD Reload utility.
11. Load DBZ, using an initial load program.
12. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9, 10, and 11 as input.
13. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 12 as input.
14. Remember to make an image copy of both databases as soon as they are loaded.

Procedure When Reorganizing DBX and DBY

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBX and DBY, using the existing DBDs and HD Unload utility.

Adding Logical Relationships

3. Code a new DBD for DBY and DBZ. “How to Specify Use of Logical Relationships in the Logical DBD” in Chapter 8, “Choosing Optional Database Functions,” on page 157 explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBX and DBY, and calculate space for DBZ.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Prereorganization utility, specifying DBR in the control statements for DBX and DBY, and DBIL in the control statement for DBZ. (The output from the Prereorganization utility says that a scan of DBX is required.)
9. Reload DBX and DBY, using the new DBDs and the HD Reload utility.
10. Load DBZ, using an initial load program.
11. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9 and 10 as input.
12. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 11 as input.
13. Remember to make an image copy of all three databases as soon as they are loaded.

Example 7. DBX and DBY Exist, DBZ Is to Be Added

Example 7 is shown in Figure 233.




Figure 233. DBX and DBY Exist, DBZ Is to Be Added

DBY must be reorganized to add the logical child pointers to the segment C prefix. Logical child pointers from segment C to segment B are not unloaded, therefore, DBX must be reorganized or scanned. DBX must be reorganized to add the logical child pointers in the segment A prefix.

Procedure Using Scan

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBY, using the existing DBD and HD Unload utility.
3. Code a new DBD for DBY and DBZ. “How to Specify Use of Logical Relationships in the Logical DBD” in Chapter 8, “Choosing Optional Database Functions,” on page 157 explains how the DBD is coded for logical relationships.

4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBY and calculate space for DBZ.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Prereorganization utility, specifying DBR in the control statements for DBY, and DBIL in the control statement for DBZ. (The output from the Prereorganization utility indicates that a scan of DBX is required.)
9. Run the scan utility against DBX.
10. Reload DBY, using the new DBDs and the HD Reload utility.
11. Load DBZ, using an initial load program.
12. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9, 10, and 11 as input.
13. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 12 as input.
14. Remember to make an image copy of both databases as soon as they are loaded.

Procedure When Reorganizing DBX and DBY

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBY and DBZ using the existing DBDs and the HD Unload utility.
3. Code a new DBD for DBY and DBZ. "How to Specify Use of Logical Relationships in the Logical DBD" in Chapter 8, "Choosing Optional Database Functions," on page 157 explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBX and DBY and calculate space for DBZ.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Prereorganization utility, specifying DBR in the control statements for DBX and DBY, and DBIL in the control statement for DBZ. (The output from the Prereorganization utility indicates that a scan of DBX is required.)
9. Reload DBX and DBY, using the new DBDs and the HD Reload utility.
10. Load DBZ, using an initial load program.
11. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9 and 10 input.
12. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 11 as input.
13. Remember to make an image copy of both databases as soon as they are loaded.

Example 8. DBX and DBY Exist, DBZ Is to Be Added

Example 8 is shown in Figure 234.




Figure 234. DBX and DBY Exist, DBZ Is to Be Added

DBY must be reorganized to add the logical child pointers in the segment C prefix. The procedure for this example (and all conditions and considerations) is exactly the same as the procedures for example 6.

Example 9. DBY Exists, DBZ Is to Be Added

Example 9 is shown in Figure 235.




Figure 235. DBY Exists, DBZ Is to Be Added

DBY must be reorganized. DBZ must be loaded using an initial load program. DBIL must be specified in the control statement for DBY. Do not specify DBR in the control statement for DBY.

Procedure

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBY, using the existing DBD and HD Unload utility.
3. Code a new DBD for DBY and DBZ. "How to Specify Use of Logical Relationships in the Logical DBD" in Chapter 8, "Choosing Optional Database Functions," on page 157 explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBY and calculate space for DBZ.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.

8. Run the Preorganization utility, specifying DBIL in the control statements for DBY and DBZ.
9. Reload DBY, using the new DBDs and the HD Reload utility.
10. Load DBZ, using an initial load program.
11. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9 and 10 as input.
12. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 11 as input.
13. Remember to make an image copy of both databases as soon as they are loaded.

Example 10. DBY Exists, DBZ Is to Be Added

Example 10 is shown in Figure 236.




Figure 236. DBY Exists, DBZ Is to Be Added

Segment X might be considered a logical child if the key of segment D is at the correct location in segment X. DBY must be reorganized, because an initial load (DBIL) of the logical parent (segment D) assumes an initial load (DBIL) of the logical child.

In this example, you could use symbolic or direct pointers for segment X. Do **not** under any circumstances specify DBR in the control statement for DBY. If you do, the reload utility will not generate work records for segment D; the logical child pointer in segment D would never be resolved. The procedure for this example (and all conditions and considerations) is exactly the same as the procedures for example 9.

Example 11. DBX and DBY Exist, DBZ Is to Be Added

Example 11 is shown in Figure 237.




Figure 237. DBX and DBY Exist, DBZ Is to Be Added

DBX and DBY must be reorganized. DBZ must be loaded using an initial load program. Because you must specify DBIL in the control statement for DBZ (a logical

Adding Logical Relationships

parent database), you must also specify DBIL for DBY (a logical child database). DBY is also a logical parent database. Therefore, you must specify DBIL in the control statement for DBX (a logical child database). The procedure for this example (and all conditions and considerations) is exactly the same as for Example 2.

Example 12. DBX and DBY Exist, DBZ Is to Be Added

Example 12 is shown in Figure 238.




Figure 238. DBX and DBY Exist, DBZ Is to Be Added

In this example, segment B has a symbolic pointer. The procedure for this example (and all conditions and considerations) is exactly the same as for example 2.

Example 13. DBX and DBY Exist, Segment Y and DBZ Are to Be Added

Example 13 is shown in Figure 239.




Figure 239. DBX and DBY Exist, Segment Y and DBZ Are to Be Added

Procedure

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload DBX, using the existing DBD and HD Unload utility.

3. Code a new DBD for DBY and DBZ. “How to Specify Use of Logical Relationships in the Logical DBD” in Chapter 8, “Choosing Optional Database Functions,” on page 157 explains how the DBD is coded for logical relationships.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for these application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space for DBX and DBY, and calculate space for DBZ.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Run the Prereorganization utility, specifying DBIL in the control statements for DBX, DBY and DBZ.
9. Reload DBX, using the new DBD and the HD Reload utility.
10. Load DBY and DBZ, using an initial load program.
11. Run the Prefix Resolution utility, using the DFSURWF1 work files that are output from Steps 9 and 10 as input.
12. Run the Prefix Update utility, using the DFSURWF3 work file that is output from Step 11 as input.
13. Remember to make an image copy of both databases as soon as they are loaded.

Steps in Reorganizing a Database to Add a Logical Relationship

Table 32 on page 447 shows you:

- When a logically related database must be scanned
- When both sides of a logical relationship must be reorganized
- When the Prefix Resolution and Prefix Update utilities must be run

The figure applies to reorganizations only. When initially loading databases, you must run the Prefix Resolution and Update utilities whenever work data sets are generated.

Table 32 covers all reorganization situations, whether or not database pointers are being changed. In using the figure, a bidirectional physically paired relationship should be treated as two unidirectional relationships. Unless otherwise specified, DBR should be specified for the reorganized databases when the Prereorganization utility is run.

The following two examples guide you in use of the figure.

Example 1. How to use Table 32

Assume your database has unidirectional symbolic pointers and you are not changing pointers. On the left side of Table 32, in the FROM column, find unidirectional symbolic pointers. The follow across to the right in the TO row and find unidirectional symbolic pointers. The figure tells you what you must do to reorganize with one of the following:

- The database containing the logical parent
- The database containing the logical child

Adding Logical Relationships

- Both databases, if necessary

In all three situations, it is not necessary to run the Prefix Resolution or Update utilities (this is what is meant by “finished”).

Example 2. How to use Table 32

Assume your database has bidirectional symbolic pointers and you need to change to bidirectional direct pointers. Table 32 shows that:

- Reorganizing only the logical parent database cannot be done, because a logical parent pointer must be created in the logical child segment in the logical child database.
- Reorganizing the logical child database can be done. To scan the logical child database, you must scan the logical parent database. The control statements for the Prereorganization utility must specify DBIL for the logical child database. Also, the Prefix Resolution and Update utilities must be run.
- Reorganizing both databases can also be done. In this case, the control statements for the Prereorganization utility must specify DBIL for the logical child database and DBR for the logical parent database. Again, the Prefix Resolution and Update utilities must be run.

Table 32. Steps in Reorganizing a Database to Add a Logical Relationship

| Type of Database | Type of Reorganization | What You Must Do to Reorganize When You Need: | | | |
|---------------------------------------|------------------------------|---|---|---|---|
| | | Unidirectional Symbolic Pointers | Unidirectional Direct Pointers | Bidirectional Symbolic Pointers | Bidirectional Direct Pointers |
| Unidirectional with symbolic pointers | Logical parent database only | Finished ¹ | Not valid, because symbolic LP pointers exist now and direct LP pointers must be added to the logical child database. | 1. Scan logical child data base. 2. Run prefix resolution and update. Note: Logical child segment will not contain LT pointers unless it is reorganized. | Not valid, because direct LP and LT pointers must be put in the logical child database. |
| | Logical child database only | Finished | 1. Scan logical parent data base. 2. Run prefix resolution and update. Specify DBIL for the logical child database. | Not valid, because a counter exists now and LCF/LCL pointers must be put into the logical parent database. | Not valid, because a counter exists now and LCF/LCL pointers must be put into the logical parent database. |
| | Both databases | Finished ² | Run prefix resolution and update. Specify DBIL for the logical child database and DBR for the logical parent database. | Run prefix resolution and update. Specify DBR for both databases. | Run prefix resolution and update. Specify DBIL for the logical child database and DBR for the logical parent database. |
| Unidirectional with direct pointers | Logical parent database only | Not valid, because a direct LP pointer exists now and symbolic LP pointers must be added to the logical child database. | 1. Scan logical child data base. 2. Run prefix resolution and update. | Not valid, because a direct LP pointer exists now and symbolic LP pointers must be added to the logical child database. LT pointers must also be added to the logical child database. | 1. Scan logical child data base. 2. Run prefix resolution and update. Note: Logical child segment will not contain LT pointers unless database is reorganized. |
| | Logical child database only | Finished | Finished | Not valid, because LCF/LCL pointers must be put in the logical parent database. | Not valid, because LCF/LCL pointers must be put in the logical parent database. |
| | Both databases | Finished ² | Run prefix resolution and update. | Run prefix resolution and update. | Run prefix resolution and update. |

Adding Logical Relationships

Table 32. Steps in Reorganizing a Database to Add a Logical Relationship (continued)

| Type of Database | Type of Reorganization | What You Must Do to Reorganize When You Need: | | | |
|--------------------------------------|------------------------------|---|--|--|---|
| | | Unidirectional Symbolic Pointers | Unidirectional Direct Pointers | Bidirectional Symbolic Pointers | Bidirectional Direct Pointers |
| Bidirectional with symbolic pointers | Logical parent database only | Not valid, because the counter in the logical parent database will not be resolved and LT pointers exist now in the logical child database. | Not valid, because symbolic LP and LT pointers exist now and a direct LP pointer must be added to the logical child database. | 1. Scan logical child data base. 2. Run prefix resolution and update. Note: LCF/LCL pointers are not unloaded and reloaded. | Not valid, because a symbolic LP pointer exists now and a direct LP pointer must be added to the logical child database. |
| | Logical child database only | Not valid, because LCF/LCL pointers exist now in the logical parent database and a counter must be added to the logical parent database. | Not valid, because LCF/LCL pointers exist now in the logical parent database and a counter must be added to the logical parent database. | 1. Scan logical parent data base. 2. Run prefix resolution and update. | 1. Scan logical parent data base. 2. Run prefix resolution and update. 3. Specify DBIL for the logical child data base. |
| | Both databases | Run prefix resolution and update. Specify DBIL for the logical child database and DBR for the logical parent database. | Run prefix resolution and update. Specify DBIL for the logical child database and DBR for the logical parent database. | Run prefix resolution and update. | Run prefix resolution and update. Specify DBIL for the logical child database and DBR for the logical parent database. |

Table 32. Steps in Reorganizing a Database to Add a Logical Relationship (continued)

| Type of Database | Type of Reorganization | What You Must Do to Reorganize When You Need: | | | |
|------------------------------------|------------------------------|--|---|--|---|
| | | Unidirectional Symbolic Pointers | Unidirectional Direct Pointers | Bidirectional Symbolic Pointers | Bidirectional Direct Pointers |
| Bidirectional with direct pointers | Logical parent database only | Not valid, because direct LP and LT pointers exist in the logical child database and symbolic LP pointers must be added. | Not valid, because the counter in the logical parent database will not be resolved and LT pointers will not be removed from the logical child database. | Not valid, because a direct LP pointer exists in the logical child database and the change is to symbolic LP pointers. | 1. Scan logical child database. 2. Run prefix resolution and update. Note: LCF/LCL pointers are not unloaded and reloaded. |
| | Logical child database only | Not valid, because LCF/LCL pointers exist in the logical parent database and a counter must be added to the logical parent database. | Not valid, because LCF/LCL pointers exist now in the logical parent database and a counter must be added to the logical parent database. | 1. Scan logical parent data base. 2. Run prefix resolution and update. | 1. Scan logical parent data base. 2. Run prefix resolution and update. |
| | Both databases | Run prefix resolution and update. Specify DBIL for the logical child database and DBR for the logical parent database. | Run prefix resolution and update. Specify DBIL for the logical child database and DBR for the logical parent database. | Run prefix resolution and update. | Run prefix resolution and update. |

Notes:

1. The Prereorganization utility says to scan the logical child database and the DFSURWF1 records will be produced if scan is run.
2. DFSURWF1 records are produced; however, the prefix resolution and update utilities need *not* be run.

Some Restrictions on Modifying Existing Logical Relationships

In some cases, the IMS utilities cannot be used to modify an existing logical relationship. When an existing logical relationship cannot be modified, you must write your own program. Two examples are as follows:

Example 1: Changing from Bidirectional Virtual to Bidirectional Physical Pairing

Figure 240 on page 450 shows the change in pairing from virtual to physical:

Adding Logical Relationships




Figure 240. The Change in Pairing from Virtual to Physical

Example 2: Changing the Location of the Real Logical Child in a Bidirectional Logical Relationship

Figure 241 shows the position change of a real logical child from one logically related database to another:




Figure 241. The Position Change of a Real Logical Child from One Logically Related Database to Another

In both of these “before” examples, occurrences of segment B can exist that are physically, but not logically, deleted. The logical child can be accessed from the logical path but not the physical path. When unloading DBX, the HD Unload utility cannot access occurrences of segment B that are physically, but not logically, deleted. Therefore, you must write your own program to do this type of reorganization.

Summary on Use of Utilities When Adding Logical Relationships

- Counters are increased by counting logical children loaded using an initial load program or, when logically related databases are reorganized, by using DBIL in the control statement.
- Counter problems can be corrected by reorganizing databases. When correcting counter problems, DBIL must be specified in the control statement for the databases involved.
- LCF and LCL pointers are not unloaded and reloaded. They must be recreated by the Prefix Resolution and Update utilities.
- Unless DBIL is specified for all its logical child databases, never specify DBIL in the control statement for a logical parent database.
- To change from symbolic to direct pointers, specify DBIL on the control statement for the logical child database.

Adding a Secondary Index

Secondary indexes are explained in Chapter 8, “Choosing Optional Database Functions,” on page 157. If you need to add a secondary index to your database:

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload your database, using the existing DBD and the HD Unload utility.
3. Code new DBDs. “How to Specify Use of Secondary Indexing in the DBD” in Chapter 6, “Choosing Full-Function Database Types,” on page 59 explains how the DBD is coded for secondary indexes. You need two new DBDs, one for the existing database and one for the new secondary index database.
4. If the change you are making affects the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Delete the old database space and define new database space (non-VSAM), or delete the space allocated for the cluster and define space for the new cluster. In addition, define space for the secondary index.
7. Reload the database, using the new DBD and the HD Reload utility.
8. Run the Prefix Resolution utility, using the DFSURWF1 work file that is output from Step 7 as input.
9. Run the HISAM unload utility, using the DFSURIDX work file that is output from Step 8 as input. Be sure to indicate in the utility control statement that HISAM unload is being used to build a secondary index.
10. Run the HISAM reload utility using as input the output from HISAM unload.
11. When you add a secondary index, remember to change your JCL. You need a DD statement for the secondary index data set even when you are not using the secondary index to process the main database. You also need to change your reorganization procedures when adding a secondary index. Whenever you reorganize the data set the secondary index points to, you must execute the reorganization utilities to rebuild the secondary index.

Adding or Converting to Variable-Length Segments

Variable-length segments are explained in Chapter 8, “Choosing Optional Database Functions,” on page 157. If you need to change selected segments in your database from fixed to variable length—or convert the entire database to variable-length segments—two ways exist to do it. Regardless of which way you use, the object in conversion is to put a size field in the segment you need to make variable length and then get the segment defined as variable length in the DBD.

Method 1. Converting Segments or a Database

To convert selected segments or the entire database this way, you must:

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Code and generate a new DBD that identifies the segment types that will be variable length, and their size.
3. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.

Adding or Converting to Variable-Length Segments

4. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
5. Write a program that sequentially retrieves from the database all segments that are to be variable length. Your program must add the 2-byte size field to each segment retrieved and then insert the segment back into the database.

Method 2. Converting Segments or a Database

To convert selected segments or the entire database this way, you must:

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload your database, using the existing DBD.
3. Code and generate a new (interim) DBD. This DBD should specify fixed-length segments for all segments being converted to variable length. It should also specify use of the segment edit/compression exit routine for each segment to be converted. (The interim DBD is used, as explained in Step 9, to add a size field to the existing fixed-length segments.)
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space if necessary. You need to do this when the change you are making results in different requirements for database space.
7. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
8. Write an edit routine to which the segment edit/compression exit routine can exit. Your edit routine should add a size field to each segment it receives. (Information on the segment edit/compression exit routine and the edit routine you must write is contained in "Segment Edit/Compression Exit Routine" on page 218.)
9. Reload the database, using the interim DBD. As each occurrence of a segment type that needs to be converted is presented for loading, your edit routine gets control and adds the size field to the segment. When your edit routine returns control, the segment is loaded into the database. Remember to make an image copy of your database as soon as it is loaded.
10. If your database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after reloading to resolve prefix information. The flowchart in Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.
11. After the database is loaded, code and generate a new DBD that specifies the segment types in the database that are variable, and their size.

Converting to the Segment Edit/Compression Exit Routine

You might need to make changes to your database before you can use a Segment Edit/Compression exit routine, such as the IMS sample exit routine DFSCMPX0.

To convert an existing database to support a Segment Edit/Compression exit routine, follow these steps:

1. Determine whether the change you are making affects the code in any application programs. If the code is affected, make sure it gets changed.
2. Unload your database, using the existing DBD and the HD Unload utility.

3. Code a new DBD. The new DBD must specify the name of your edit routine for the segment types you need edited.
4. If the change you are making affected the code in application programs, make any necessary changes to the PSBs for those application programs. If you have the DB/DC Data Dictionary, it can help you determine which application programs and PCBs are affected by the DBD changes you have made.
5. Rebuild the ACB if you have ACBs prebuilt rather than built dynamically.
6. Recalculate database space. You need to do this because the change you are making results in different requirements for database space.
7. Delete the old database space and define new database space. If you are using VSAM, use the Access Method Services DEFINE CLUSTER command to define VSAM data sets.
8. Reload the database, using the new DBD. Remember to make an image copy of your database as soon as it is reloaded.
9. If your database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after reloading to resolve prefix information. Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.

Related Reading: For more information on Segment Edit/Compression exit routines, see:

- “Segment Edit/Compression Exit Routine” on page 218
- *IMS Version 9: Customization Guide*

Converting Databases for Data Capture Exit Routines and Asynchronous Data Capture

This topic contains general-use programming interface information.

Data Capture exit routines are explained in “Data Capture Exit Routines” on page 221. To convert an existing database for use with Data Capture exit routines or Asynchronous Data Capture:

1. Determine whether the change requires revisions to the logical delete rules in a database. If so, change the delete rules, which might require reorganizing your database.
2. Code a new DBD. On the DBD or SEGM statements, specify the name of each exit routine you need called against a segment in the database.

Related Reading:

- See *IMS Version 9: Utilities Reference: System* for details on the DBD parameters required for Data Capture exit routines or Asynchronous Data Capture.
 - *IMS Version 9: Customization Guide* explains the exit routines in detail, how to code them, and how they work.
3. Run DBDGEN.
 4. If you use prebuilt ACBs rather than dynamically built ACBs, rebuild the ACB.

Converting a Logical Parent Concatenated Key from Virtual to Physical or Physical to Virtual

You can convert a logical parent concatenated key from virtual to physical or from physical to virtual by using DBDGEN and the HD reorganization utilities. To do this conversion:

1. Unload your database, using the existing DBD.
2. Code a new DBD, changing the concatenated key physical/virtual specification.
3. If you use prebuilt ACBs rather than dynamically built ACBs, rebuild the ACB.
4. Recalculate the database space. You need to do this because the change you are making changes database space requirements.
5. For non-VSAM data sets, delete the old database space and define new database space. For VSAM data sets, delete the space allocated for the old clusters and define space for the new clusters.
6. If your database uses logical relationships or secondary indexes, you must run some of the reorganization utilities before and after reloading to resolve prefix information. Figure 197 on page 352 tells you which utilities to use and the order in which they must be run.
7. Reload your database using the new DBD. Remember to make an image copy of your database as soon as it is reloaded.
8. If required, run reorganization utilities to resolve prefix information.

Using the Online Change Function

Using the online change function, you can add, change, and delete full-function databases, HALDB master databases, and DEDB databases online without stopping IMS. MSDBs do not support the online change function.

The online change function for DEDBs allows both database-level and area-level changes. A database-level change affects the structure of the DEDB and includes such changes as adding or deleting an area, adding a segment type, or changing the randomizing routines. An area-level change involves increasing or decreasing the size of an area (IOVF, DOVF, CI). An area-level change requires the user to stop only that area with the /DBRECOVERY command; a database-level change requires the user to stop all areas of the DEDB.

Unlike standard randomizers which distribute database records across the entire DEDB, two-stage randomizers distribute database records within an area. By using a two-stage randomizer, changes to an individual area's root addressable allocation are area-level changes, and only the areas affected need to be stopped.

To use online change with full-function databases, HALDB master databases, and DEDB databases, perform the following steps:

1. Allocate the required new data sets (see *IMS Version 9: Installation Volume 1: Installation Verification* for planning these data sets).
2. Run a MODBLKS system definition if additions, changes, or deletions to the system definition DATABASE (and possibly APPLCTN) statements need to be made (see *IMS Version 9: Administration Guide: System* for more information).
3. Run the necessary DBDGEN (see *IMS Version 9: Utilities Reference: Database and Transaction Manager*), PSBGEN, and ACBGEN (see *IMS Version 9: Utilities Reference: System*).

Note: All changes to ACBLIB members resulting from the ACBGEN execution are available to the online system after the online change (provided that the changed resources—PSBs and DBDs—are defined in the online system).

4. Update the security definitions of the IMS system's security facilities to include any new databases. Security facilities can include RACF, another external security product, the IMS Security Maintenance utility, and exit routines. For more information on IMS security, see *IMS Version 9: Administration Guide: System*.
5. Allocate the database data sets for databases to be added.
6. Load your database.
7. For Fast Path, online change must be completed before the database can be loaded. Also, Fast Path can only load databases online; batch jobs cannot be used.
8. If dynamic allocation is used in a z/OS environment, run the dynamic allocation utility.
9. Use the online change utility to copy your updated staging libraries to the inactive libraries (see *IMS Version 9: Utilities Reference: System* for information on running this utility).
10. Issue the operator commands to cause your inactive libraries to become your active libraries (see *IMS Version 9: Command Reference* for the commands used).

If a database in a z/OS environment needs to be reorganized because of changes to the active ACBLIB data set, the /DBR command must be issued to deallocate the database prior to the /MODIFY COMMIT command that introduces the ACBGEN changes. The commands /DBR, /DBD, or /STA DATABASE ACCESS= must be completed to take the areas of the database to be changed or deleted offline prior to issuing the /MODIFY COMMIT command.

Restriction: For HALDB databases, if you delete a database name from IMS using online change, you cannot reuse that name for another HALDB partition without first performing a cold start. Similarly, if you delete a partition name from DBRC, you cannot use online change to reuse that name for another database without first performing a cold start.

Online Change, DEDBs, and Availability of IFP and MPP Regions

Changes can be made to DEDBs using online change while maintaining the availability of IFP and MPP regions that access the DEDBs. If database level changes are made to the DEDB while an IFP/MPP is running, then the application will pseudoabend and the PSB will be rescheduled on the next DL/I call to the DEDB.

Two level changes can be made to DEDBs. The database level changes allow:

1. Add or Delete DEDBs.
2. Add or Delete segment types.
3. Add, Change, or Delete a segment and its fields.
4. Add, Change, or Delete segment compression routines.
5. Add, Change, or Delete data capture exit routines.
6. Change randomizers.
7. Add or Delete areas.

Using the Online Change Function

8. Change area RAP space allocation and the randomizer is not a 2-stage randomizer.

The area level changes allow:

1. Change area RAP space allocation and the randomizer is a 2-stage randomizer.
2. Change DOVF or IOVF space allocation.
3. Change SDEP space allocation.
4. Change CI size.

Area level changes and items 4 through 8 of the database level change require a BUILD DBD (not a BUILD PSB). In this case, with exception to items 4 and 5 when the defined PSB SENSEGs have reference to exit routines that are added or deleted, the PSB does not change. Changes can be made to DEDBs using online change while maintaining the availability of IFP and MPP regions that access the DEDBs only if there is no change to the scheduled PSB. The application will then pseudoabend with ABENDU0777 and the PSB will be rescheduled on the next DL/I call to DEDB. The message DFS2834I is issued. Other changes to the PSBs such as items 1 through 5 of the DEDB database changes, full-function database changes, or PSB changes using online change require that the IFP and MPP regions be brought down.

The following procedure describes the steps necessary to make database level changes to a DEDB with an IFP / MPP running:

1. Use a specific user-developed application program or OEM utility to unload the DEDB through existing system definitions.
2. DBDGEN, PSBGEN and ACBGEN to generate the application control blocks to implement the DEDB structural changes. The changed or new application control blocks must be built into the active IMS system's staging copy of ACBLIB, which is offline.
3. Run the online change utility, DFSUOCU0, to move the changed ACBLIB from the staging ACBLIB to the inactive (A or B) copy of the ACBLIB that is online to the active IMS system.
4. Enter the normal /DBR command sequence to remove access to the DEDB from the active IMS system.
5. Enter and follow the online change command sequence for PREPARE processing for ACBLIB changes.
6. Enter and follow the online change command sequence for COMMIT/ABORT processing for ACBLIB changes. The online IMS system will switch from using the active (A or B) copy of the ACBLIB to the inactive (A or B) copy.
7. Delete, define and initialize all of the DEDB AREA data sets with the new system definitions.
8. Enter the normal /START DATABASE and /START AREA commands to make the DEDB and its AREAs accessible to the active IMS system.
9. Use a specific user-developed application program or OEM utility to reload the DEDB through the change system definitions for the DEDB.
10. On the first access to the newly changed DEDB, the application will pseudoabend and the PSB will be rescheduled. Message DFS2834I will be displayed.

The transaction will be tried again for both IFPs and MPPs when the PSB is rescheduled. If the application attempts to access the DEDB before commit

processing has completed, an 'FH' status will be returned to the application. The DEDB is inaccessible because the randomizer for the DEDB is unloaded by the /DBR command.

If either database level or area level changes are made to DEDBs while a BMP or DBCTL thread is active, then commit processing fails and the message DFS3452 is issued.

Related Reading: See the *IMS Version 9: Messages and Codes, Volume 2* for more information on message DFS3452 and other messages.

Online Change and DEDB Randomizer and Exit Routines

Randomizing routines determine the location of database records by AREA within the DEDB and by root anchor point (RAP) within the AREA. A change of the DEDB randomizer is a database level change. A new randomizing routine affects the location (AREA and RAP) of every database record within the DEDB. The randomizer is defined for the DEDB in the DBD parameter: RMNAME=.

A randomizer change can involve introducing a brand new randomizer into the active IMS system or changing an existing randomizer in use by one or more DEDBs.

Introducing a New DEDB Randomizing Routine Using Online Change

The name of the randomizer is specified in the DBD parameter: RMNAME=. If a new randomizer is introduced for an existing DEDB, a DBDGEN and ACBGEN of the database with the new randomizer name is required in addition to the following procedural steps:

1. Use a specific customer-developed application program or original equipment manufacturer (OEM) utility to unload the DEDB with the current randomizer.
2. Assemble and link edit the new randomizer into the IMS SDFSRESL or one of the libraries in the IMS SDFSRESL STEPLIB concatenation.
3. Run a DBDGEN for the DEDB with the new randomizer designated in the DBD parameter: RMNAME=.
4. ACBGEN is also needed to build the application control blocks to implement the DEDB definition that includes the new randomizer. The changed or new application control blocks must be built into the active IMS system's staging copy of ACBLIB, which is offline.
5. ACBLIB Run the online change utility, DFSUOCU0, to move the changed ACBLIB from the staging ACBLIB to the inactive (A or B) copy of the ACBLIB that is online to the active IMS system.
6. Enter the normal /DBR operator command sequence to remove access to the DEDB from the active IMS system.
7. Enter and follow the online change command sequence for PREPARE processing for ACBLIB changes.
8. Enter and follow the online change command sequence for COMMIT/ABORT processing for ACBLIB changes. The online IMS system will switch from using the active (A or B) copy of the ACBLIB to the inactive (A or B) copy.
9. Delete, define and initialize all of the DEDB AREA data sets with the new system definitions.
10. Enter the normal /START DATABASE and /START AREA commands to make the DEDB and its areas accessible to the active IMS system.

Using the Online Change Function

11. Use a specific customer-developed application program or OEM utility to reload the DEDB with the new randomizing routine in effect.

Changing DEDB Randomizing Routines Using Online Change

If a change is made to a randomizer already in use by one or more DEDBs, then all of the DEDBs using the subject randomizer must be included in the change process.

The changed randomizer will not be introduced if an existing version is already loaded for any DEDB in the active IMS system. You can determine that the existing version is no longer used by locating the keyword GONE in message DFS2838I. Also, you can determine that the randomizer module is brought from any library to the storage by locating the keyword LOADED in the message DFS2842I.

Changing DEDB randomizers requires the procedures described below. Because the name of the randomizer remains the same, DBDGEN, ACBGEN and the online change command sequence are not applicable.

1. Use a specific customer-developed application program or OEM utility to unload the DEDB with the existing randomizer. This should be done for all of the DEDBs that use the randomizer to be changed.
2. Enter the normal /DBR DATABASE operator command sequence to remove access to the DEDBs from the active IMS system. The /DBR DATABASE command unloads the randomizer for the DEDBs designated as operands. When all the DEDBs that reference the randomizer are stopped, the randomizer is removed from the active IMS system. If a DEDB is not stopped and references a randomizer that has been removed from the IMS system, then a U1021 abend results on the next DL/I call.
3. Assemble and link edit the changed randomizer into the IMS SDFSRESL or one of the libraries of the IMS SDFSRESL STEPLIB concatenation.
4. Delete, define and initialize all of the DEDB AREA data sets to prepare for reloading the DEDB with the changed randomizer.
5. Enter the /START DATABASE command for each of the DEDBs that use the changed randomizer. For DEDBs, the /START DATABASE command causes the randomizer to be loaded.
6. Use a specific customer-developed application program or OEM utility to reload the DEDB with the changed randomizing routine in effect.

Deleting a DEDB Randomizing Routine Using Online Change

To delete a randomizer from the active IMS system, follow the procedural steps that are documented under “Introducing a New DEDB Randomizing Routine Using Online Change” on page 457. Once all the DEDBs that were using the old randomizer have been unloaded and had the /DBR command run successfully against them, then the randomizer can be deleted. Customers with data sharing IMS systems that do not share SDFSRESLs must be careful to delete the randomizer from both systems. A message (DFS2838) is generated when the randomizer is deleted.

Online Change and DEDB Segment Compression Routines

Segment compression routines are segment specific and are defined for the DEDB in the DBD SEGM parameter (“COMPRTN=”). Adding, changing or deleting segment compression routines is procedurally the same and involves the same restrictions as DEDB randomizing routines.

Online Change and DEDB Data Capture Exit Routines

Data Capture exit routines can be defined for the DEDB on the DBD statement, for a specific segment on the SEGM statement ("EXIT="), or for both. Multiple exit routines can be specified on a single DBD or SEGM statement.

Adding a New DEDB Data Capture Exit Routine Using Online Change: To add a new Data Capture exit routine, follow the procedure below:

1. Assemble and link edit the new exit routine into the IMS.SDFSRESL or one of the libraries in the IMS.SDFSRESL STEPLIB concatenation.
2. Run a DBDGEN for the DEDB with the new exit routine designated in the DBD or SEGM parameter: "EXIT=".
3. ACBGEN is also needed to build the application control blocks to implement the DEDB definition that includes the new exit routine. The changed or new application control blocks must be built into the active IMS system's staging copy of ACBLIB, which is offline.
4. Run the online change Utility, DFSUOCU0, to move the changed ACBLIB from the staging ACBLIB to the inactive (A or B) copy of the ACBLIB that is online to the active IMS system.
5. Enter the normal /DBR command sequence to remove access to the DEDB from the active IMS system.
6. Enter and follow the online change command sequence for PREPARE processing for ACBLIB changes.
7. Enter and follow the online change command sequence for COMMIT/ABORT processing for ACBLIB changes. The online IMS system will switch from using the active (A or B) copy of the ACBLIB to the inactive (A or B) copy.
8. Enter the normal /START DATABASE and /START AREA commands to make the DEDB and its areas accessible to the active IMS system.

Changing an Existing DEDB Data Capture Exit Routine Using Online Change:

To change an existing Data Capture exit routine, follow these steps:

1. Allow the dependent regions that are accessing DEDBs with the particular Data Capture exit to end normally.
2. Assemble and link edit the changed exit routine into the IMS SDFSRESL or one of the libraries of the IMS SDFSRESL STEPLIB concatenation.
3. Start the dependent regions. Data Capture exits are loaded at dependent region initialization time, so the new version of the exit will take effect when the region is started. Data Capture exit routines that were linked as reentrant or reusable are unloaded at dependent region termination time. Otherwise, they are unloaded after every DL/I call.

Deleting a DEDB Data Capture Exit Routine Using Online Change: To delete a Data Capture exit routine, execute the following steps:

1. Run a DBDGEN for the DEDB with the old exit routine omitted from the DBD or SEGM statement.
2. Build the application control blocks by using the Application Control Block Maintenance utility to implement the DEDB definition that excludes the old exit routine. The changed or new application control blocks must be built into the active IMS system's staging copy of ACBLIB, which is offline.
3. Run the Online Change Copy utility, DFSUOCU0, to move the changed ACBLIB from the staging ACBLIB to the inactive (A or B) copy of the ACBLIB that is online to the active IMS system.

Using the Online Change Function

4. Enter the normal /DBR command sequence to remove access to the DEDB from the active IMS system.
5. Enter and follow the online change command sequence for PREPARE processing for ACBLIB changes.
6. Enter and follow the online change command sequence for COMMIT/ABORT processing for ACBLIB changes. The online IMS system will switch from using the active (A or B) copy of the ACBLIB to the inactive (A or B) copy.
7. Enter the normal /START DATABASE and /START AREA commands to make the DEDB and its areas accessible to the active IMS system.

Changing DEDB Root Addressable Space with Two Stage Randomizer Using Online Change

The UOW structure and root addressable allocation is specific to each area within each DEDB. However, a change to the number of root addressable CIs within one area can affect the number of root anchor points within the whole DEDB. If the DEDB uses a standard randomizing routine that randomly distributes database records across the entire database, changes to the root addressable allocation are *Database Level* changes and procedurally must be handled as such. This topic is not applicable to such changes.

If, however, a two-stage randomizer is used for the DEDB, a change to an individual area UOW root addressable definition is an *AREA Level* change. A two-stage randomizer does not attempt to evenly distribute database records across all areas based on the total number of root anchor points in the entire DEDB. A two-stage randomizer is designated in the DBDGEN by coding the randomizer name as follows:

RMNAME=(mmmmmmmm,2)

In prior releases of IMS, customers would get the following error message if a DEDB DBD had more than one operand in the RMNAME parameter:

8, DBD130 - RMNAME OPERAND IS OMITTED OR INVALID

The same message will appear for this release of IMS if anything but a two is specified as the second operand of RMNAME. Customers can still specify RMNAME=(mmmmmmmm) for standard randomizing routines.

Changing the DEDB AREA UOW Structural Definition Using Online Change

Changing the DEDB AREA UOW structural definition requires the following procedural steps:

1. Use a specific customer-developed application program or original equipment manufacturer (OEM) utility to unload the area through existing system definitions.
2. DBDGEN, PSBGEN and ACBGEN to generate the application control blocks to implement the DEDB structural changes. The "UOW=(x,y)" parameter on the AREA DBDGEN macro statement defines the amount of space allocated to overflow within a DEDB UOW. The "ROOT=(nnn,mmm)" parameter on the AREA DBDGEN macro statement defines the amount of space allocated to Independent Overflow.

The changed or new application control blocks must be built into the active IMS system's staging copy of ACBLIB, which is offline.

3. Run the online change utility, DFSUOCU0, to move the changed ACBLIB from the staging ACBLIB to the inactive (A or B) copy of the ACBLIB that is online to the active IMS system.

4. Enter the /DBR AREA command to remove access to the area from the active IMS system.
5. Enter and follow the online change command sequence for PREPARE processing for ACBLIB changes.
6. Enter and follow the online change command sequence for COMMIT/ABORT processing for ACBLIB changes.
7. Delete, define and initialize the area with the new system definitions.
8. Enter the /START AREA command to make the area accessible to the active IMS system.
9. Use a specific customer-developed application program or OEM utility to reload the DEDB through the changed system definitions for the DEDB.

Making Online Changes at the DEDB and Area Level

This topic contains the following information about making online changes to DEDB and DEDB areas:

- “Adding or Deleting DEDBs Using Online Change”
- “Adding or Deleting DEDB Segments Using Online Change” on page 462
- “Adding or Deleting DEDB Areas Using Online Change” on page 463
- “Changing DEDB Root Addressable Space Allocation Using Online Change” on page 463
- “Online Change and DEDB Dependent and Independent Overflow Space Allocation” on page 464
- “Changing DEDB CI Size Using Online Change” on page 464

Adding or Deleting DEDBs Using Online Change

Figure 242 shows the overall process for adding a database using online change.

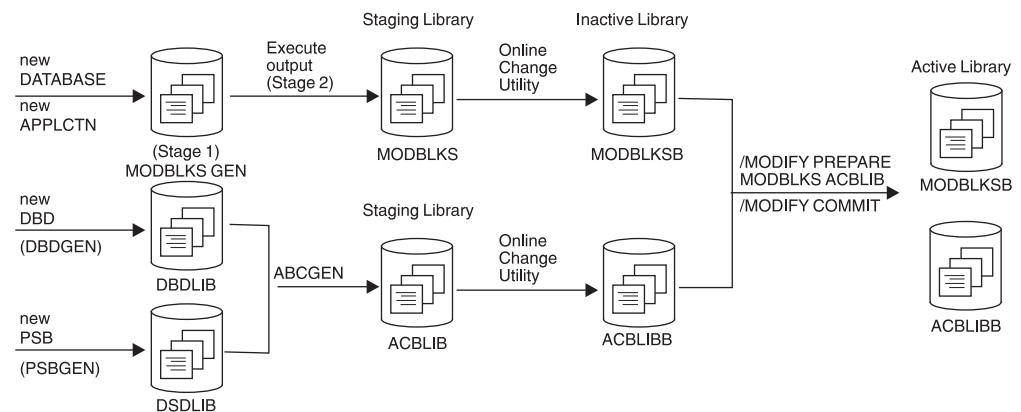


Figure 242. Adding a Database Using Online Change

Adding or deleting a DEDB and implementing the change by means of the IMS online change facility requires that you follow the steps described below. See Figure 242 for an overall picture.

1. MODBLKs Level system definition (Stage 1 and Stage 2) to add or delete the DEDB. The changed MODBLKs should be generated into the active IMS system's staging copy of MODBLKs, which is offline.
2. DBDGEN, PSBGEN and ABCGEN to generate the application control blocks to add or delete the DEDB and PSBs that access it. The changed or new

Using the Online Change Function

application control blocks must be generated into the active IMS system's staging copy of ACBLIB, which is offline.

3. Run the online change utility, DFSUOCU0, to move the changed MODBLKS and ACBLIB changes from the staging libraries to the inactive (A or B) copies of these libraries that are online to the active IMS system.
4. Enter and follow the online change command sequence for PREPARE processing. If a DEDB is being added to an IMS system that does not have Fast Path installed, the DFS2833 error message will appear and the PREPARE process will be aborted.

If a DEDB is added whose areas have CI sizes that exceed the system buffer size (BSIZ=), then message DFS2832 will appear and the PREPARE process aborts.

Finally, if a DEDB is added to an IMS system that was initialized without any DEDBs, then message DFS2837 appears and the PREPARE process aborts.

Output threads are initialized during Fast Path initialization only if DEDBs are currently generated in the system. In order for the user to be able to add DEDBs with online change, IMS must be initialized with DEDBs to begin with.

5. If the DEDB is to be deleted, any BMP region or DBCTL thread scheduled for access to the DEDB must first be stopped. Full function transactions scheduled for access to the DEDB will be placed in a QSTOP state and as a result, MPP or IFP dependent regions need not be stopped to implement the online change to delete the DEDB.
6. If the DEDB is to be deleted, access to it from the active IMS system must be removed by means of a /DBR DB command. The commit will fail with a DFS3452 message if the DEDB has not had the /DBR command successfully run against it beforehand.
7. Execute the online change command sequence for COMMIT/ABORT processing.
8. If the DEDB is newly added, execute the following additional steps at any appropriate time prior to making the DEDB generally available for normal user access:
 - a. Execute the normal procedures for defining the new DEDB and its areas and area data sets to DBRC and the RECON data sets.
 - b. Define and initialize all of the area data sets belonging to the new DEDB.
 - c. Execute the procedures to include the required Dynamic Allocation definitions that will enable the DEDB and its areas to be allocated to the active IMS system. Or register the DEDB and its areas to DBRC, and DBRC will dynamically allocate them during IMS initialization.
 - d. Enter the /START DATABASE and /START AREA commands to make the DEDB and its areas accessible to the active IMS system.
 - e. Run the necessary application load programs.

Related Reading: See the *IMS Version 9: Messages and Codes, Volume 2* for information on the types of messages you might receive while adding or deleting DEDBs.

Adding or Deleting DEDB Segments Using Online Change

Adding or deleting segment types or changing segment formats affects the structure of a DEDB and constitutes a Database Level change. The addition or deletion of segment types (including the DEDB Sequential Dependent Segment type) affects the hierarchical structure and the segment prefix layout to implement this structure. Similarly, the change of individual segment formats changes the structure of the entire database and space allocations within each AREA of the DEDB.

To make structural changes to an existing DEDB, execute the procedural steps described below.

1. Use a specific customer-developed application program or OEM utility to unload the DEDB through existing system definitions.
2. DBDGEN, PSBGEN and ACBGEN to generate the application control blocks to implement the DEDB structural changes. The changed or new application control blocks must be built into the active IMS system staging copy of ACBLIB, which is offline.
3. Run the online change utility, DFSUOCU0, to move the changed ACBLIB from the staging ACBLIB to the inactive (A or B) copy of the ACBLIB that is online to the active IMS system.
4. Enter the normal /DBR command sequence to remove access to the DEDB from the active IMS system. This command may be issued any time prior to the /MODIFY COMMIT.
5. Enter and follow the online change command sequence for PREPARE processing for ACBLIB changes.
6. Enter and follow the online change command sequence for COMMIT/ABORT processing for ACBLIB changes.
7. Delete, define and initialize all of the AREA data sets belonging to the DEDB with the new system definitions.
8. Enter the normal /START DATABASE and /START AREA commands to make the DEDB and its areas accessible to the active IMS system.
9. Use a specific customer-developed application program or OEM utility to reload the DEDB through the changed system definitions for the DEDB.

Adding or Deleting DEDB Areas Using Online Change

Adding or deleting an area can affect the location of every database record throughout the DEDB. Changing the number of areas will alter the number of root anchor points (RAPs) within the DEDB. DEDB randomizing routines attempt to randomly distribute database records throughout the entire DEDB based first on the area and then on the root anchor point (RAP) within the area.

Adding or deleting one or more areas to a DEDB constitutes a structural change such as adding a segment type. The steps described in “Adding or Deleting DEDB Segments Using Online Change” on page 462 should be followed to change the number of areas defined in the DEDB. If areas are newly added, the required DBRC definitions for areas and area data sets must be processed and dynamic allocation blocks must be prepared before the new areas can be accessed by the active IMS system.

Changing DEDB Root Addressable Space Allocation Using Online Change

There are different implications depending on whether you randomly distribute DEDB records or use a standard randomizer to evenly distribute DEDB records. In either case, you can distribute DEDB records across an entire DEDB or just a single DEDB area.

Random Distribution of Database Records Across All DEDB Areas: Changes to the DEDB unit of work (UOW) structure that affect the number of DEDB control intervals defined to the Root Addressable portion impact the number of root anchor points within the entire DEDB. This type of change potentially affects the location of every database record within the DEDB.

Using the Online Change Function

Standard DEDB Randomizing Routines: Standard DEDB randomizing routines attempt to evenly distribute database records across all areas and within the selected AREA. Such randomizers determine the record location based on the total number of root anchor points in the entire DEDB.

A change to the UOW structure that changes the number of CIs defined to the root addressable area constitutes Database Level change when a standard DEDB randomizing routine is used. This type of change should be treated the same as a DEDB structural change in terms of online change procedures.

Online Change and DEDB Dependent and Independent Overflow Space Allocation

Starting in IMS Version 3, Fast Path has provided limited support for extending DEDB AREA Independent Overflow space allocation. That support continues unchanged. Additionally, DEDB online change will allow changes to the overflow space allocation both within each UOW (Dependent Overflow) and outside the root addressable portion (Independent Overflow) of the AREA. Both Dependent and Independent Overflow changes are considered to be Area-level changes. However, such changes must not alter the number of CIs defined to the root addressable portion. Changing the number of root addressable CIs will change the number of root anchor points and could affect the DEDB randomizing routine in locating database records.

Changing DEDB AREA overflow allocation requires the same procedural steps as those defined for changing the root addressable area.

Related Reading: See “Changing the DEDB AREA UOW Structural Definition Using Online Change” on page 460 for details on changing the DEDB AREA overflow.

Changing DEDB CI Size Using Online Change

DEDB online change can be used to change DEDB AREA control interval size. However, CI size changes must not alter the number of CIs allocated to the root addressable portion of an AREA because this could affect the DEDB randomizer in locating database records across the DEDB. The SIZE= parameter on the AREA statement of DBDGEN defines the CI size of the data set that constitutes the AREA.

Extending DEDB Independent Overflow Online

You can extend the independent overflow (IOVF) portion of a DEDB area while IMS is online by following the procedure described in this topic. The first time the area is opened after this procedure is completed, a message is issued to verify that Fast Path recognizes and accepts the change to the area and normal open processing completes. You can also modify the IOVF portion of a DEDB using DEDB online change.

You cannot decrease the size of the IOVF with this procedure. However, the size of the sequential dependent part might increase or decrease depending on the total amount of space allocated to the area. The steps in this procedure also reorganize the area.

To increase the size of the IOVF portion of a DEDB online you must:

1. Run the DBDGEN utility to obtain an updated DBD. Update *only* the following operands on the ROOT= keyword of the AREA statement:

number

Specifies the total number of units of work (UOWs) allocated to the root addressable and the IOVF parts of the area. Increase **number** to reflect the number of UOWs you need to add to the IOVF.

overflow

Specifies the space reserved for the IOVF, expressed as the number of UOWs. Increase the number on this operand by the same amount you increase the **number** operand. For example, if the original values were **number**= x and **overflow**= y , and if **number** is changed to $x + 2$, then **overflow** must be changed to $y + 2$.

All other control statements must remain identical to those on the existing DBD. Changing other control statements might damage data and create unpredictable results.

2. Run the ACBGEN utility using the updated DBD. You should run PSB=ALL to create a new and complete ACBLIB with the new ROOT= parameters. The output should be a different data set from the one currently used by the control region. The new ACBLIB is identical to the old ACBLIB, except for the ROOT= changes. You can use the staging ACBLIB, but do not switch with the online change function.
3. Ensure that the area is in good condition. The area must not have any in-doubts, and must not be in a recovery-needed condition. Also, at least one copy of the area (one area data set) must have no error queue elements (EQEs). Use the /DIS AREA command to display EQEs and the condition. Use the /DIS CCTL INDOUBT command to display all in-doubt threads. Eliminate potential defects before continuing to the next step so that data is not lost or damaged.
4. Process SDEPs using the SDEP scan and delete utilities. This step is required because the IOVF extension procedure requires an unload and load of the area. Some unload and load utilities are unable to process SDEPs. Unload/load utilities that do process SDEPs might reload them in root order rather than time order, which can interfere with subsequent SDEP scan and delete operations.

Related Reading:

- For more information on the DBRC definitions for the shared AREAs with SDEP segments, see the *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.
 - For more information on DEDB Sequential Dependent Scan utility keywords and change boundaries, see the *IMS Version 9: Utilities Reference: Database and Transaction Manager*.
 - For more information on the DEDB Sequential Dependent Scan utility user-written exit routine parameter interface, see the *IMS Version 9: Customization Guide*.
5. If multiple copies of the area (MADS) exist, stop all copies of the area except one using the /STOP ADS command. Ensure that the remaining copy does not have any EQEs and is not in a recovery-needed condition. Multiple ADSs must be stopped to ensure that DBRC has accurate information when the area is brought online after the IOVF is extended.
 6. Issue a /DBR or /STO AREA command against the area.
 7. Take an image copy of the area.
 8. If the area is registered with DBRC, set the recovery-needed flag on for the area. This flag is required by the DEDB Initialization utility and can be set using a CHANGE.DBDS RECOV command.

Extending DEDB Independent Overflow Online

9. Unload the area.
10. Execute the IDCAMS utility to delete and redefine the data set. The amount of space you allocate for the area in the Define procedure should reflect the increased size of the IOVF. The number of SDEP CIs in the area might change because this number represents the difference between the total amount of space allocated to the area and the amount used by the other parts. These other parts are the root addressable part, the IOVF, the reorganization UOW, and two control CIs.

Related Reading: See *DFSMS Access Method Services for Catalogs* for a description of the IDCAMS Delete and Define functions.

11. Execute the Fast Path initialization utility against the new area using the new ACBLIB.
12. Issue the /START AREA command to bring the area online.
13. Reload the area.

Note: It is recommended that you reload the area in batch. If you reload the area using a BMP, the BMP might fail with message DFS3709A and reason code 5. If this failure occurs, issue the CHANGE.DBDS command to set ICOFF and restart the BMP.

IMS Version 9: Messages and Codes, Volume 2 explains message DFS3709A and the reason for this failure.

14. Take an image copy of the area after the reload.

When the area is next accessed, message DFS3703I is issued. This message alerts you that discrepancies were found during open processing. However, open processing continues because the discrepancies indicate to IMS that you used an accepted procedure to increase the size of the IOVF. DFS3703I is not issued during subsequent opens of the area as long as IMS remains online. DFS3703I is also issued by any sharing subsystem the first time the area is opened on that subsystem after the IOVF is extended.

During emergency restart or extended recovery facility (XRF) takeover, the updated area information is picked up from the log. Therefore, DFS3703I is not issued.

Use the new ACBLIB for any subsequent normal restarts of the online system. Ensure that the new ACBLIB reflects *only* the changes made to the ROOT= keyword. Any other changes you make might cause damage to the area. If you do not use the new ACBLIB, open logic allows the discrepancy between information from the old ACBLIB and information from the area data set, but issues message DFS3703I each time the discrepancy is encountered.

Note: Remember that you cannot use the online change function to update the ACBLIB with the altered ROOT= parameter.

Part 3. Appendixes

Appendix A. Meaning of Bits in the Delete Byte

This appendix examines the meanings of:

- “Bits in the Delete Byte”
- “Bits in the Prefix Descriptor Byte”

Bits in the Delete Byte

This topic contains diagnosis, modification or tuning information.

The meaning of each bit in the delete byte, when turned on, is as follows:

| Bit | Meaning When Delete Byte is Turned On |
|-----|--|
| 0 | Segment has been marked for deletion. This bit is used for segments in a HISAM or secondary index database or segments in primary index. |
| 1 | Database record has been marked for deletion. This bit is used for segments in a HISAM or secondary index database or segments in a primary index. |
| 2 | Segment has been processed by the delete routine. |
| 3 | This bit is reserved. |
| 4 | Prefix and data portion of the segment are separated in storage. (The delete byte preceding the separated data portion of the segment has all bits turned on.) |
| 5 | Segment has been marked for deletion from a physical path. This bit is called the PD (physical delete) bit. |
| 6 | Segment has been marked for deletion from a logical path. This bit is called the LD (logical delete) bit. |
| 7 | Segment has been marked for removal from its logical twin chain. This bit should only be set on if bits 5 and 6 are also on). |

Bits in the Prefix Descriptor Byte

This topic contains diagnosis, modification, or tuning information.

The delete byte is also used for the root segment of a DEDB, only there it is called a prefix descriptor byte. The meaning of each bit, when turned on, is as follows:

| Bit | Meaning When Root Segment Prefix Descriptor is Turned On |
|-----|--|
| 0 | Sequential dependent segment is defined. |
| 1-3 | These bits are reserved. |
| 4-7 | If the number of defined segments is 8 or less, bits 4 through 7 contain the highest defined segment code. Otherwise, the bits are set to 000. |

Appendix B, “Insert, Delete, and Replace Rules for Logical Relationships,” on page 471, discusses replacing, inserting, and deleting rules for logical relationships, which includes how to specify rules in a physical DBD and a rules summary.

Bits in the Prefix Descriptor Byte

Appendix B. Insert, Delete, and Replace Rules for Logical Relationships

You need to examine all your application requirements and decide who can insert, delete, and replace segments involved in logical relationships and how those updates are to be made (physical path only or physical and logical path). The insert, delete, and replace rules in the physical DBD determine how updates apply across logical relationships.

This appendix examines the following information on rules:

- “Specifying Rules in the Physical DBD”
- “Insert Rules” on page 472
- “Replace Rules” on page 475
- “Using the DLET Call” on page 481

This appendix contains general-use programming interface information.

Specifying Rules in the Physical DBD

Insert, delete, and replace rules are specified using the RULES= keyword of a SEGM statement in the DBD for logical relationships. Figure 243 contains a diagram of the RULES= keyword and its parameters.




Figure 243. Insert, Delete, and Replace Rules in the DBD

The valid parameter values for the RULES= keyword are:

- B** Specifies a bidirectional virtual delete rule. It is not a valid value for either the first or last positional parameter of the RULES= keyword.
- L** Specifies a logical insert, delete, or replace rule.
- P** Specifies a physical insert, delete, or replace rule.
- V** Specifies a virtual insert, delete, or replace rule.

The RULES= keyword accepts three positional parameters:

- The first positional parameter sets the insert rule
- The second positional parameter sets the delete rule
- The third positional parameter sets the insert rule

For example, RULES=P,L,V says the insert rule is physical, the delete rule is logical, and the replace rule is virtual. The B rule is only applicable for delete. In general, the P rule is the most restrictive, the V rule is least restrictive, and the L rule is somewhere in between.

The RULES= parameter is applicable only to segments involved in logical paths, that is, the logical child, logical parent, and physical parent segments. The RULES= parameter is not coded for the virtual logical child.

Insert Rules

The insert rules apply to the destination parent segments, but not to the logical child segment. A destination parent can be a logical or physical parent. The insert rule has no meaning for the logical child segment except to satisfy the RULES= macro's coding scheme. Therefore, any insert rule (P, L, V) can be coded for a logical child. A logical child can be inserted provided:

- The insert rule of the destination parent is not violated
- The logical child being inserted does not already exist (it cannot be a duplicate)

A description of how the insert rules work for the destination parent is as follows:

- When RULES=P is specified, the destination parent can be inserted *only* using the physical path. This means the destination parent must exist before inserting a logical path. A concatenated segment is not needed, and the logical child is inserted by itself. Figure 244 on page 473 shows an example of the physical insert rule.
- When RULES=L is specified, the destination parent can be inserted either using the physical path or concatenated with the logical child and using the logical path. When a logical child/destination parent concatenated segment is inserted, the destination parent is inserted if it does not already exist and the I/O area key check does not fail. If the destination parent does exist, it will remain unchanged and the logical child will be connected to it. Figure 247 on page 474 shows an example of the logical insert rule.
- When RULES=V is specified, the destination parent can be inserted either using the physical path or concatenated with the logical child and using the logical path. When a logical child/destination parent concatenated segment is inserted, the destination parent is replaced if it already exists. If it does not already exist, the destination parent is inserted. Figure 249 on page 474 shows an example of the virtual insert rule.

The Logical Child Insert Call

To insert the logical child segment, the I/O area in an application program must contain one of the following segments in accordance with the destination parent's insert rule:

- The logical child
- The logical child/destination parent concatenated segment

For all DL/I calls, either an error is detected and an error status code returned (in which case no data is changed), or the required changes are made to all segments effected by the call. Therefore, if the required function cannot be performed for both parts of the concatenated segment, an error status code is returned, and no change is made to either the logical child or the destination parent.

The insert operation is not affected by KEY or DATA sensitivity as specified in a logical DBD or a PCB. This means that if a program is other than DATA sensitive to both the logical child and destination parent of a concatenated segment, and if the insert rule is L or V, the program must still supply both of them in the I/O area when inserting using a logical path. Because of this, maintenance programs that insert concatenated segments should be DATA sensitive to both segments in the concatenation.

Status Codes

The nonblank status codes that can be returned to an application program after an ISRT call are as follows:

- AM—An insert was attempted and PROCOPTI
- GE—The parent of the destination parent or logical child was not found
- II—An attempt was made to insert a duplicate segment
- IX—The rule specified was P, but the destination parent was not found

One reason for getting an IX status code is that the I/O area key check failed. Concatenated segments must contain the destination parent's key twice—once as part of the logical child's LPCK and once as a field in the parent. These keys must be equal.

Figure 244, Figure 245, and Figure 246 show a physical insert rule example.

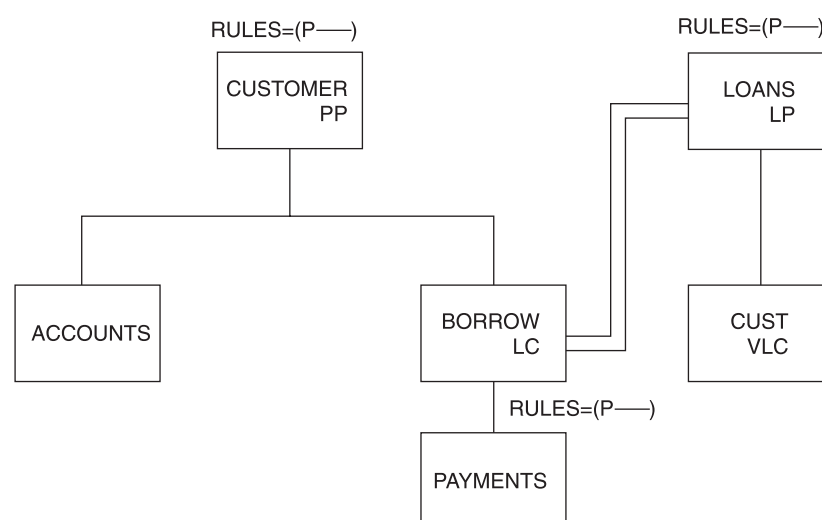


Figure 244. Physical Insert Rule Example

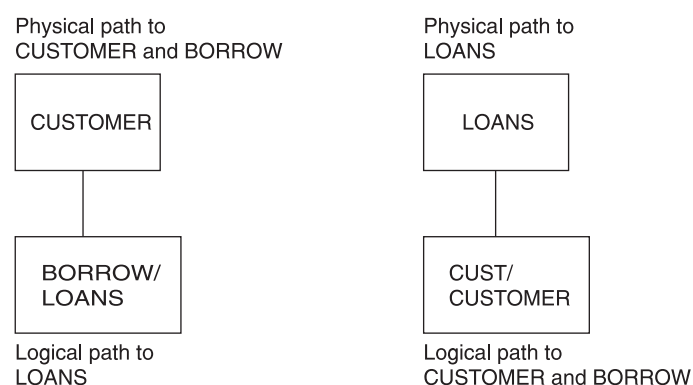


Figure 245. Paths for Physical Insert Rule Example

```

ISRT 'CUSTOMER' STATUS CODE=' '
ISRT 'BORROW' STATUS CODE=' ' ('IX' if LOANS does not exist)
  
```

Figure 246. ISRT and Status Codes for Physical Insert Rule Example

Insert Rules

Figure 247 and Figure 248 show a logical insert rule example.

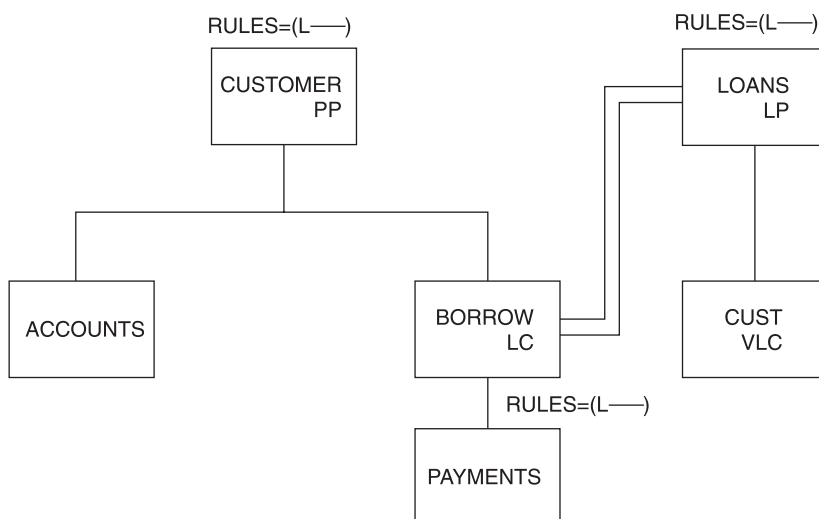


Figure 247. Logical Insert Rule Example

```
ISRT 'LOANS' STATUS CODE=' '  
ISRT 'CUST' STATUS CODE='IX'
```

Figure 248. ISRT and Status Codes for Logical Insert Rule Example

The IX status code shown in Figure 248 is the result of omitting the concatenated segment CUST/CUSTOMER in the second call. IMS checked for the key of the CUSTOMER segment (in the I/O area) and failed to find it. With the L insert rule, the concatenated segment must be inserted to create a logical path.

Figure 249 and Figure 250 show a virtual insert rule example.

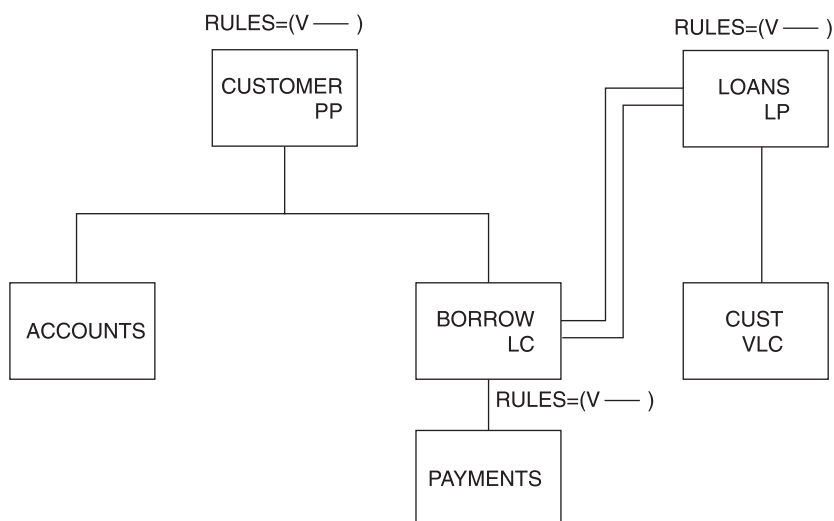


Figure 249. Virtual Insert Rule Example

```
ISRT 'CUSTOMER' STATUS CODE=' '  
ISRT 'BORROW/LOANS' STATUS CODE=' '
```

Figure 250. ISRT and Status Codes for Virtual Insert Rule Example

The code shown in Figure 250 will replace the LOANS segment if present, and insert the LOANS segment if not. The V insert rule is a powerful option.

Insert Rules Summary

Specifying the insert rule as P prevents inserting the destination parent as part of a concatenated segment. A destination parent can only be inserted using the physical path. If the insert creates a logical path, only the logical child needs to be inserted.

Specifying the insert rule as L on the logical and physical parent allows insertion using either the physical path or the logical path as part of a concatenated segment. When inserting a concatenated segment, if the destination parent already exists it remains unchanged and the logical child is connected to it. If the destination parent does not exist, it is inserted. In either case, the logical child is inserted if it is not a duplicate, and the destination parent's insert rule is not violated.

The V insert rule is the most powerful of the three rules. The V insert rule is the most powerful because it will insert the destination parent (inserted as a concatenated segment using the logical path) if the parent did not previously exist, or otherwise replace the existing destination parent with the inserted destination parent.

Replace Rules

The replace rules are applicable to the physical parent, logical parent, and logical child segments of a logical path. The following is a description of how the replace rules work:

- When RULES=P is specified, the segment can only be replaced when retrieved using a physical path. If this rule is violated, no data is replaced and an RX status code is returned. Figure 251 on page 476 shows an example of the physical replace rule.
- When RULE=L is specified, the segment can only be replaced when retrieved using a physical path. If this rule is violated, no data is replaced. However, no RX status code is returned, and a blank status code is returned. Figure 253 on page 477 shows an example of the logical replace rule.
- When RULES=V is specified, the segment can be replaced when retrieved by either a physical or logical path. Figure 255 on page 478 shows an example of the virtual replace rule.

The Replace Call

A replace operation can be done only on that portion of a concatenated segment to which an application program is data sensitive. If no data is changed in a segment, no data is replaced. Therefore, no replace rule is violated. The replace rule is not checked for a segment that is part of a concatenated segment but is not retrieved.

For all DL/I calls, either an error is detected and an error status code returned (in which case no data is changed), or the required changes are made to all segments affected by the call. Therefore, if the required function cannot be performed for both parts of the concatenated segment, an error status code is returned, and no change is made to either the logical child or the destination parent.

Replace Rule Status Codes

The status code returned to an application program indicates the first violation of the replace rule that was detected. These status codes are as follows:

Replace Rules

- AM—a replace was attempted and PROCOPTR
- DA—the key field of a segment or a non-replaceable field was changed
- RX—the replace rule was violated

Figure 251 and Figure 252 show a physical replace rule example.

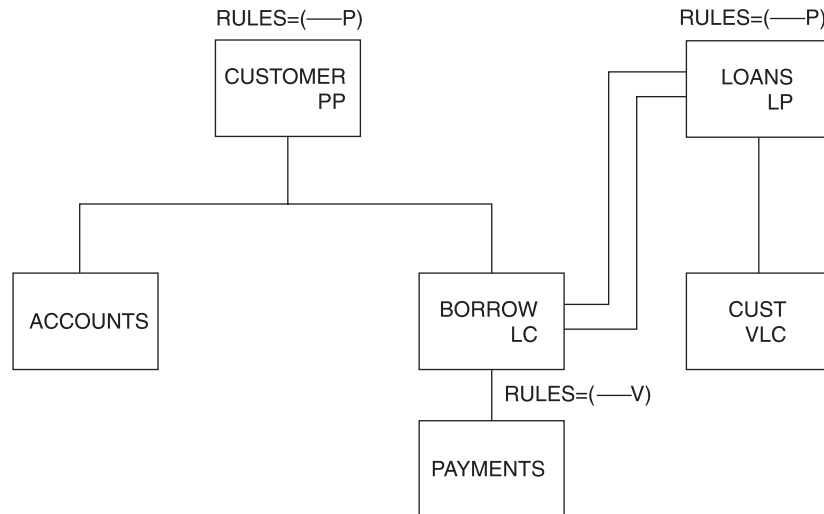


Figure 251. Physical Replace Rule Example

```
GHU 'CUSTOMER' STATUS CODE=' '
REPL          STATUS CODE=' '
GHN 'BORROW/LOANS' STATUS CODE=' '
REPL          STATUS CODE='RX'
```

Figure 252. Calls and Status Codes for Physical Replace Rule Example

The P replace rule prevents replacing the LOANS segment as part of a concatenated segment. Replacement must be done using the segment's physical path.

Figure 253 on page 477 and Figure 254 on page 477 show a logical replace rule example.

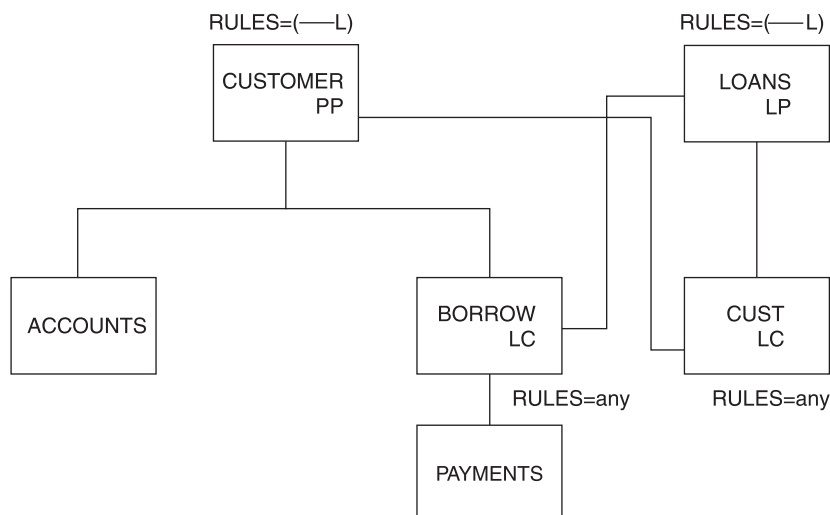


Figure 253. Logical Replace Rule Example

```

GHU 'CUSTOMER'
    'BORROW/LOANS' STATUS CODE=' '
REPL STATUS CODE=' '
  
```

Figure 254. Calls and Status Codes for Logical Replace Rule Example

As shown in Figure 253, the L replace rule prevents replacing the LOANS segment as part of a concatenated segment. Replacement must be done using the segment's physical path. However, the status code returned is blank. The BORROW segment, accessed by its physical path, is replaced. Because the logical child is accessed by its physical path, it does not matter which replace rule is selected.

The L replace rule allows replacing only the logical child half of the concatenation, and the return of a blank status code.

Figure 255 on page 478 and Figure 256 on page 478 show a virtual replace rule example.

Replace Rules

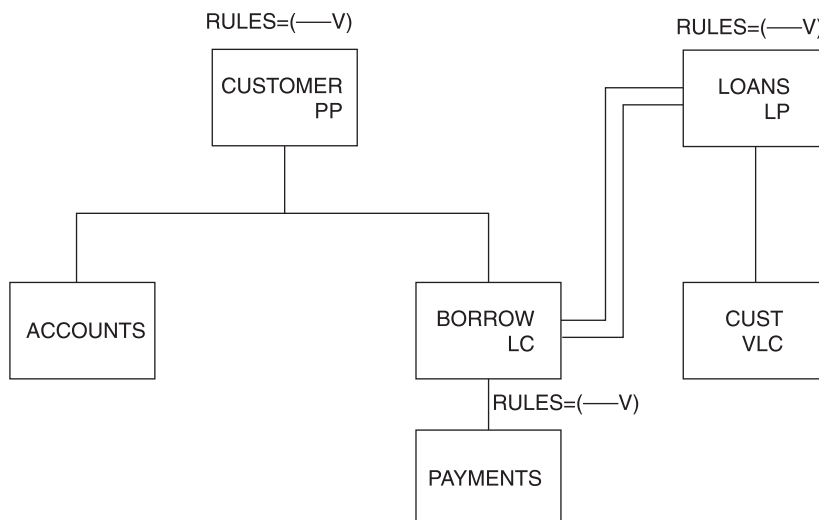


Figure 255. Virtual Replace Rule Example

```
GHU 'LOANS'
    'CUST/CUSTOMER' STATUS CODE=' '
REPL STATUS CODE=' '
```

Figure 256. Calls and Status Codes for Virtual Replace Rule Example

As shown in Figure 256, the V replace rule allows replacing the CUSTOMER segment using its logical path as part of a concatenated segment.

Replace Rules Summary

Specifying the replace rule as P, on any segment in a logical relationship, prevents replacing that segment except when it is retrieved using its physical path. When the replace rule for the logical parent is specified as L, IMS returns a blank status code without replacing any data when the logical parent is accessed concatenated with the logical child. Because the logical child has been accessed by its physical path, its replace rule can be any of the three. So, using the replace rule allows the selective replacement of the logical child half of the concatenation and a blank status code. Specifying a replace rule of V, on any segment of a logical relationship, allows replacing that segment by either its physical or logical path.

Table 33 on page 479 and Table 34 on page 480 show all of the possible combinations of replace rules that can be specified. They show what actions take place for each combination when a call is issued to replace a concatenated segment in a logical database. Table 33 on page 479 and Table 34 on page 480 are based on the databases and logical views shown in Figure 257 on page 479 and Figure 258 on page 479.

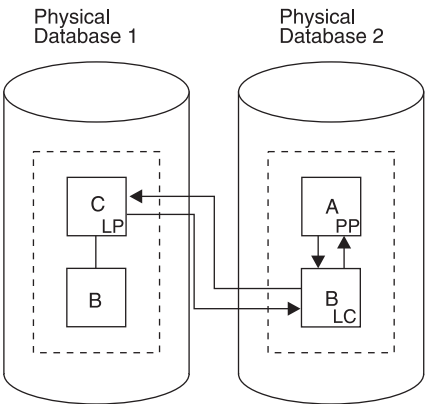


Figure 257. Physical Databases for Replace Rules Tables




Figure 258. Logical Views for Replace Rules Table

Table 33. Replace Rules for Logical View 1

| Replace Rule Specified | | Segment Attempting to Replace | | Status Code | Data Replaced? | |
|------------------------|---|-------------------------------|---|-------------|----------------|---|
| | | B | C | | B | C |
| P | P | X | | | Y | |
| P | P | | X | RX | | N |
| P | P | X | X | RX | N | N |
| P | L | X | | | Y | |
| P | L | | X | | | N |
| P | L | X | X | | Y | N |
| P | V | X | | | Y | |
| P | V | | X | | | Y |
| P | V | X | X | | Y | Y |
| L | P | X | | | Y | |
| L | P | | X | RX | | N |
| L | P | X | X | RX | N | N |
| L | L | X | | | Y | |
| L | L | | X | | | N |
| L | L | X | X | | Y | N |
| L | V | X | | | Y | |
| L | V | | X | | | Y |

Replace Rules

Table 33. Replace Rules for Logical View 1 (continued)

| Replace Rule Specified | | Segment Attempting to Replace | | Status Code | Data Replaced? | |
|------------------------|---|-------------------------------|---|-------------|----------------|---|
| | | B | C | | B | C |
| L | V | X | X | | Y | Y |
| V | P | X | | | Y | |
| V | P | | X | RX | | N |
| V | P | X | X | RX | N | N |
| V | L | X | | | Y | |
| V | L | | X | | | N |
| V | L | X | X | | Y | N |
| V | V | X | | RX | Y | |
| V | V | | X | RX | | Y |
| V | V | X | X | RX | Y | Y |

Table 34. Replace Rules for Logical View 2

| Replace Rule Specified | | Segment Attempting to Replace | | Status Code | Data Replaced? | |
|------------------------|---|-------------------------------|---|-------------|----------------|---|
| | | B | A | | B | A |
| P | P | X | | RX | Y | |
| P | P | | X | | | N |
| P | P | X | X | RX | N | N |
| P | L | X | | RX | Y | |
| P | L | | X | | | N |
| P | L | X | X | RX | Y | N |
| P | V | X | | RX | Y | |
| P | V | | X | | | Y |
| P | V | X | X | RX | Y | Y |
| L | P | X | | | Y | |
| L | P | | X | RX | | N |
| L | P | X | X | RX | N | N |
| L | L | X | | | Y | |
| L | L | | X | | | N |
| L | L | X | X | | Y | N |
| L | V | X | | | Y | |
| L | V | | X | | | Y |
| L | V | X | X | | Y | Y |
| V | P | X | | RX | Y | |
| V | P | | X | RX | | N |
| V | P | X | X | | N | N |
| V | L | X | | | Y | |
| V | L | | X | | | N |
| V | L | X | X | | Y | N |

Table 34. Replace Rules for Logical View 2 (continued)

| Replace Rule Specified | | Segment Attempting to Replace | | Status Code | Data Replaced? | |
|------------------------|---|-------------------------------|---|-------------|----------------|---|
| | | B | A | | B | A |
| V | V | X | | | Y | |
| V | V | | X | | | Y |
| V | V | X | X | | Y | Y |

Using the DLET Call

The DLET call is a request to delete a path of segments, not a request to release the DASD space used by a segment. Delete rules are needed when a segment is involved in a logical relationship, because that segment belongs to two paths: a physical and a logical path. The selection of the delete rules for the logical child and its logical and physical parent (or two logical parents if physical pairing is used) determines whether one or two DLET calls are necessary to delete the two access paths.

Physical and Logical Deletion

Physically deleting a segment prevents further access to that segment using its physical parents. Physically deleting a segment also physically deletes its physical dependents, however one exception to this exists: If one of the physical parents of the physically deleted segment is a logical child that has been accessed from its logical parent, then the physically deleted segment is accessible from that logical child. The deleted segment is accessible from that logical child because the physical dependents of a logical child are variable intersection data.

Logically deleting a logical child prevents further access to the logical child using its logical parent. Unidirectional logical child segments are assumed to be logically deleted. A logical parent is considered logically deleted when all its logical children are physically deleted. For physically paired logical relationships, the physical child paired to the logical child must also be physically deleted before the logical parent is considered logically deleted.

Deleting Concatenated Segments

The following application program can be sensitive to either the concatenated segment—SOURCE=(DATA/DATA), (DATA/KEY), (KEY/DATA)—or the logical child, because it is the logical child that is either physically or logically deleted (depending on the path accessed) in all cases. The concatenated segment relationships are shown in Figure 259 on page 482.

Delete Rules

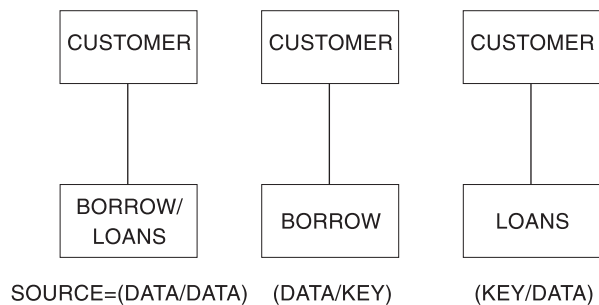


Figure 259. Concatenated Segment Relationships

The Third Access Path

In Figure 260, three paths to the logical child segment SEG4 exist:

- The physical path from its physical parent SEG3
- The logical path from its logical parent SEG7
- A third path from SEG4's physical dependents (SEG5 and SEG6) (because segment SEG6 is a logical parent accessible from its logical child SEG2)

Related Reading: See “Possibility of Abnormal Termination” on page 503 for more information on potential abends.

These paths are called “full-duplex” paths, which means accessibility to segments in the paths is in two directions (up and down). Two delete bits that control access along the paths exist, but they are “half-duplex,” which means they only block half of each respective path. No bit that blocks the third path exists. If SEG4 were both physically and logically deleted (in which case the PD and LD bits in SEG4 would be set), SEG4 would still be accessible from the third path, and so would both of its parents.

Neither physical nor logical deletion prevents access to a segment from its physical or logical children. Logically deleting SEG4 prevents access to SEG4 from its logical parent SEG7, and it does not prevent access from SEG4 to SEG7. Physically deleting SEG4 prevents access to SEG4 from its physical parent SEG3, but it does not prevent access from SEG4 to SEG3.

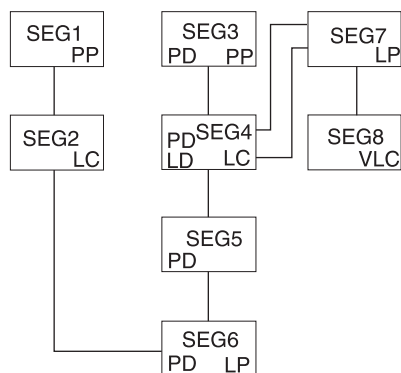


Figure 260. Third Access Path Example

Use of the Delete Byte

The delete byte is used by IMS to maintain the delete status of segments within a database. The meaning of each bit within the delete byte is in “Bits in the Delete Byte” on page 469. The bit is only meaningful for logical child segments and their logical parents. For segments involved in a logical relationship, the PD and LD bits are set or assumed set as follows:

- If a segment is physically deleted (thereby preventing further access to it from its physical parent), then delete processing scans downward from the deleted segment through its dependents, turns upward, and either releases each segment’s DASD space or sets the PD bit. HISAM is the one exception to this process. In HISAM, the delete bit is set in the segment specified by the DLET call and processing terminates.
- If the PD bit is set in a logical parent, the LD bit is set in all logical children that can be reached from that logical parent.
- When physical pairing is used, if the PD bit is set in one of a pair of logical children, the LD bit is set in its paired segment.
- When a virtually paired logical child is logically deleted (thereby preventing further access to it from its logical parent), the LD bit is set in the logical child.
- The LD bit is assumed set in all logical children in unidirectional logical relationships.
- If physical pairing is used, the LD bit is assumed set in a parent if all the paired segments that are physical children of the parent have the PD bit set on.

Issuing the Delete Call

A DLET call can be issued against a segment defined in either a physical or logical DBD. The call can be issued against either a physical segment or a concatenated segment.

A DLET call issued against a concatenated segment requests deletion of the logical child in the path that is accessed. If a concatenated segment or a logical child is accessed from its logical parent, the DLET call requests logical deletion. In all other cases, a delete call requests physical deletion.

Physical deletion of a segment generates a request for logical deletion of all the segment’s logical children and generates a request for physical deletion of all the segment’s physical children. Physical deletion of a segment also generates a request to delete any index pointer segments for which the physically deleted segment is the source segment.

Delete sensitivity must be specified in the PCB for each segment against which a delete call can be issued. The call does not need to be specified for the physical dependents of those segments. Delete operations are not affected by KEY or DATA sensitivity as specified in either the PCB or logical DBD.

Status Codes

The nonblank status codes that can be returned to an application program after a DLET call are as follows:

- DX—A delete rule was violated
- DA—The key was changed in the I/O area
- AM—The call function was not compatible with the processing option or segment sensitivity

Delete Rules

DASD Space Release

The DLET call is not a request for release of DASD space. Depending on the database organization, DASD space can or cannot be reused when it is released. DASD space for a segment is released when the following conditions are met:

- Space has been released for all physical dependents of the segment.
- The segment is physically deleted (PD bit is set or being set on).
- If the segment is a logical child or logical parent, then it must be physically and logically deleted (PD bit is set or being set on and LD bit is set or assumed set).
- If the segment is a dependent of a logical child (and is variable intersection data) and the DLET call was issued against a physical parent of the logical child, the logical child must be both physically and logically deleted.
- If the segment is a secondary index pointer segment, the space has been released for its target segment.

Delete Rules

The following is a description of how the delete values work for the logical parent, physical parent, and logical child.

Logical Parent

- When RULES=P is specified, the logical parent must be logically deleted before a DLET call is effective against it or any of its physical parents. Otherwise, the call results in a DX status code, and no segments are deleted. However, if a delete request is made against a segment as a result of propagation across a logical relationship, then the P rule acts like the L rule that follows.
- When RULES=L is specified, either physical or logical deletion can occur first. When the logical parent is processed by a DLET call, all logical children are logically deleted, but the logical parent remains accessible from its logical children.
- When RULES=V is specified, a logical parent is deleted along its physical path explicitly when deleted by a DLET call. All of its logical children are logically deleted, although the logical parent remains accessible from these logical children.

A logical parent is deleted along its physical path implicitly when it is no longer involved in a logical relationship. A logical parent is no longer involved in a logical relationship when:

- It has no logical children pointing to it (its logical child counter is zero, if it has any)
- It points to no logical children (all of its logical child pointers are zero, if it has any)
- It has no physical children that are also real logical children

Physical Parent (Virtual Pairing Only)

- PHYSICAL/LOGICAL/VIRTUAL is meaningless.
- BIDIRECTIONAL VIRTUAL means a physical parent is automatically deleted along its physical path when it is no longer involved in a logical relationship. A physical parent is no longer involved in a logical relationship when:
 - It has no logical children pointing to it (its logical child counter is zero, if it has one)
 - It points to no logical children (all of its logical child pointers are zero, if it has any)
 - It has no physical children that are also real logical children

Logical Child

- When RULES=P is specified, the logical child segment must be logically deleted first and physically deleted second. If physical deletion is attempted first, the DLET call issued against the segment or any of its physical parents results in a DX status code, and no segments are deleted. If a delete request is made against the segment as a result of propagation across a logical relationship, or if the segment is one of a physically paired set, then the rule acts like the L rule that follows.
- When RULES=L is specified, deletion of a logical child is effective for the path for which the delete was requested. Physical and logical deletion of the logical child can be performed in any order. The logical child and any physical dependents remain accessible from the non-deleted path.
- When RULES=V is specified, a logical child is both logically and physically deleted when it is deleted through either its logical or physical path (setting either the PD or LD bits sets both bits). If this rule is coded on only one logical child segment of a physically paired set, it acts like the L rule.

Note: For logical children involved in unidirectional logical relationships, the meaning of all three rules is the same, so any of the three rules can be specified.

Examples Using the Delete Rules

Figure 261 through Figure 296 show the use of the delete rules for each of the segment types for which the delete rule can be coded (logical and physical parents and their logical children). Only the rule pertinent to the example is shown in each figure. The explanation accompanying the example applies only to the specific example.

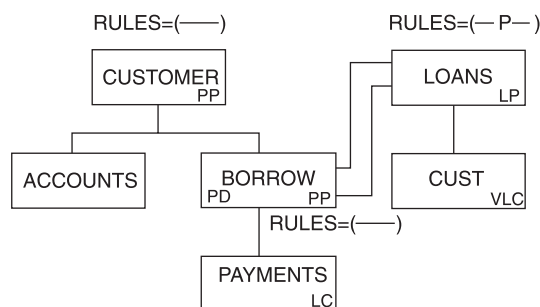


Figure 261. Logical Parent, Virtual Pairing—Physical Delete Rule Example

Delete Rules




Figure 262. Logical Parent, Physical Pairing—Physical Delete Rule Example: Before and After

```
GHU 'LOANS' STATUS=' '
DLET STATUS=' '

```

Figure 263. Logical Parent, Physical Pairing—Physical Delete Rule Example: Database Calls

The physical delete rule requires that all logical children be previously physically deleted. Physical dependents of the logical parent are physically deleted.

The DLET status code will be 'DX' if all of the logical children were not previously physically deleted. All logical children are logically deleted. The LD bit is set on in the physical logical child BORROW.

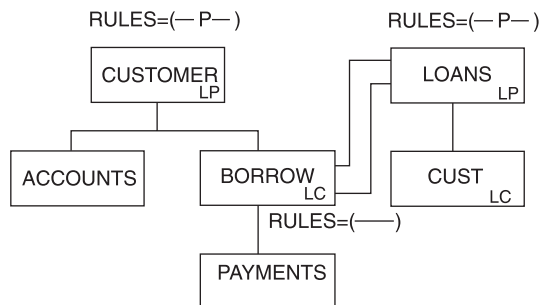


Figure 264. Logical Parent, Physical Pairing—Physical Delete Rule Example




Figure 265. Logical Parent, Physical Pairing—Physical Delete Rule Example: Before and After

```
GHU 'CUSTOMER' STATUS=' '
DLET           STATUS=' '  '
```

Figure 266. Logical Parent, Physical Pairing—Physical Delete Rule Example: Calls and Status Codes

The physical delete rule requires that:

- All logical children be previously physically deleted.
- Physical children paired to the logical child be previously deleted.

CUSTOMER, the logical parent, has been physically deleted. Both the logical child and its pair had previously been physically deleted. (The PD and LD bits are set on the before figure of the BORROW/LOANS.)

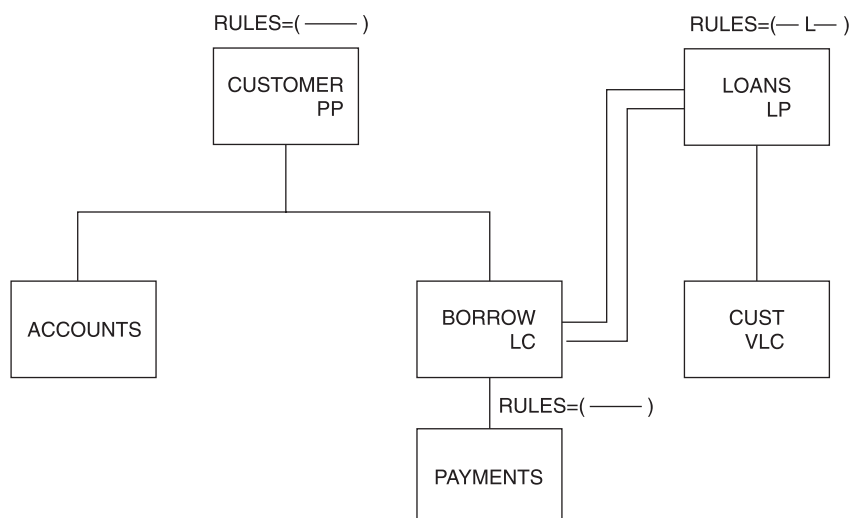


Figure 267. Logical Parent, Virtual Pairing—Logical Delete Rule Example

Delete Rules




Figure 268. Logical Parent, Virtual Pairing—Logical Delete Rule Example: Before and After

```
GHU 'LOANS' STATUS=' '
DLET      STATUS=' ';
```

Figure 269. Logical Parent, Virtual Pairing—Logical Delete Rule Example: Calls and Status Codes

The logical delete rule allows either physical or logical deletion first; neither causes the other. Physical dependents of the logical parent are physically deleted.

The logical parent LOANS remains accessible from its logical children. All logical children are logically deleted. The LD bit is set on in the physical child BORROW.

The processing and results shown in Figure 267 on page 487 would be the same if the logical parent LOANS delete rule were virtual instead of logical. The example that follows is an additional one to explain the logical delete rule.

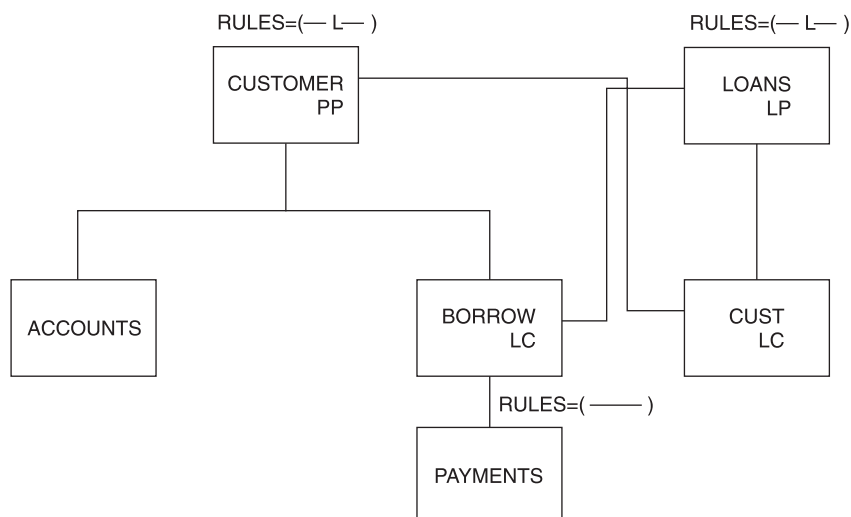


Figure 270. Logical Parent, Physical Pairing—Logical Delete Rule Example




Figure 271. Logical Parent, Physical Pairing—Logical Delete Rule Example: Before and After

```

GHU 'LOANS' STATUS=' '
DLET STATUS=' '
  
```

Figure 272. Logical Parent, Physical Pairing—Logical Delete Rule Example: Calls and Status Codes

The logical delete rule allows either physical or logical deletion first; neither causes the other. Physical dependents of the logical parent are physically deleted.

The logical parent LOANS remains accessible from its logical children. All physical children are physically deleted. Paired logical children are logically deleted.

Delete Rules

The processing and results shown in Figure 270 on page 489 would be the same if the logical parent LOANS delete rule were virtual instead of logical. An additional example to explain the virtual delete rule follows in Figure 273.

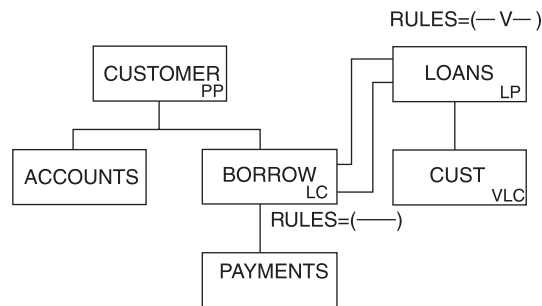


Figure 273. Logical Parent, Virtual Pairing—Virtual Delete Rule Example




Figure 274. Logical Parent, Virtual Pairing—Virtual Delete Rule Example: Before and After

```
GHU 'CUSTOMER'
    'BORROW/LOANS' STATUS=' '
DLET STATUS=' '

```

Figure 275. Logical Parent, Virtual Pairing—Virtual Delete Rule Example: Calls and Status Codes

The virtual delete rule allows explicit and implicit deletion. Explicit deletion is the same as using the logical rule. Implicit deletion causes the logical parent to be physically deleted when the last logical child is physically deleted.

Physical dependents of the logical child are physically deleted. The logical parent is physically deleted. Physical dependents of the logical parent are physically deleted. The LD bit is set on in the physical logical child BORROW.

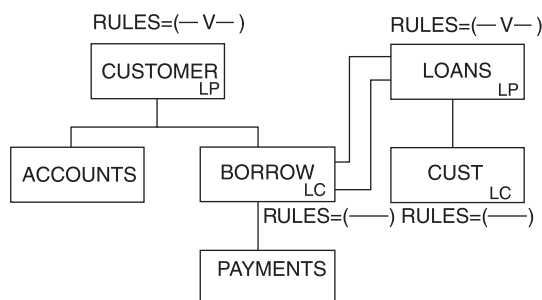



Figure 276. Logical Parent, Physical Pairing—Virtual Delete Rule Example

Before



After




Figure 277. Logical Parent, Physical Pairing—Virtual Delete Rule Example: Before and After

```

GHU 'CUSTOMER'
    'BORROW/LOANS' STATUS=' '
DLET                                STATUS=' '

```

Figure 278. Logical Parent, Physical Pairing—Virtual Delete Rule Example: Calls and Status

The virtual delete rule allows explicit and implicit deletion. Explicit deletion is the same as using the logical rule. Implicit deletion causes the logical parent to be physically deleted when the last logical child is physically and logically deleted.

The logical parent is physically deleted. Any physical dependents of the logical parent are physically deleted.

Note: The CUST segment must be physically deleted before the DLET call is issued. The LD bit is set on in the BORROW segment.

Delete Rules

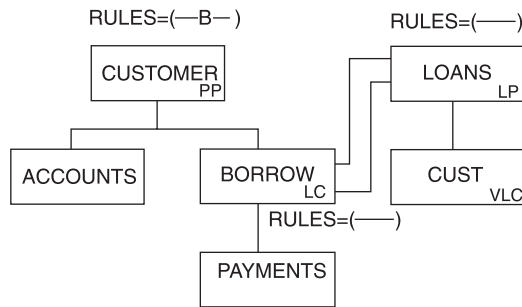


Figure 279. Physical Parent, Virtual Pairing—Bidirectional Virtual Example




Figure 280. Physical Parent, Virtual Pairing—Bidirectional Virtual Example: Before and After

```

GHU 'LOANS'
    'CUSTOMER' STATUS=' '
DLET          STATUS=' '
  
```

Figure 281. Deleting Last Logical Child Deletes Physical Parent

The bidirectional virtual rule for the physical parent has the same effect as the virtual rule for the logical parent.

When the last logical child is logically deleted, the physical parent is physically deleted. The logical child (as a dependent of the physical parent) is physically deleted. All physical dependents of the physical parent are physically deleted. That is, ACCOUNTS (not shown), BORROW, and PAYMENT are physically deleted.

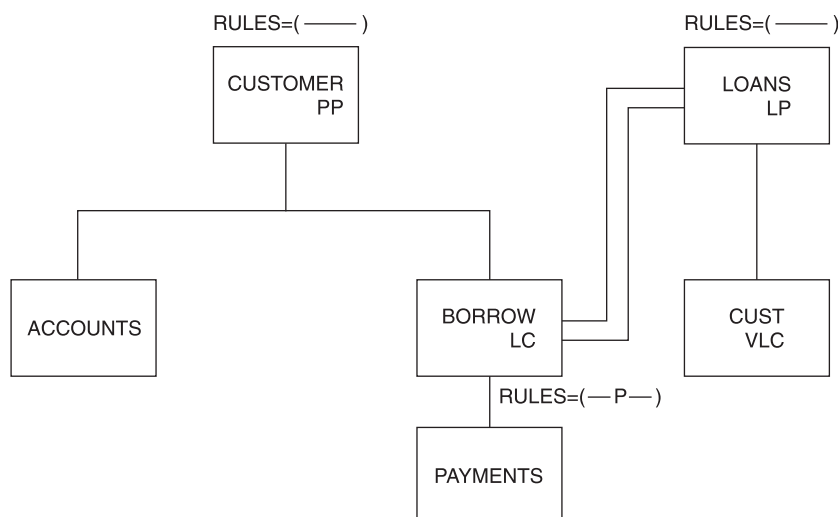


Figure 282. Logical Child, Virtual Pairing—Physical Delete Rule Example

```

GHU 'LOANS'          STATUS=' '
    'CUST/CUSTOMER'
DLET STATUS=' '

GHU 'CUSTOMER'       STATUS=' '
    'BORROW/LOANS'
DLET STATUS=' '
  
```

Figure 283. Logical Child, Virtual Pairing—Physical Delete Rule Example: Deleting the Logical Child

The physical delete rule requires that the logical child be logically deleted first. The LD bit is now set in the BORROW segment.

The logical child can be physically deleted only after being logically deleted. After the second delete, the LD and PD bits are both set. The physical delete of the logical child also physically deleted the physical dependents of the logical child. The PD bit is set.

Delete Rules




Figure 284. Logical Child, Virtual Pairing—Physical Delete Rule Example: Before and After

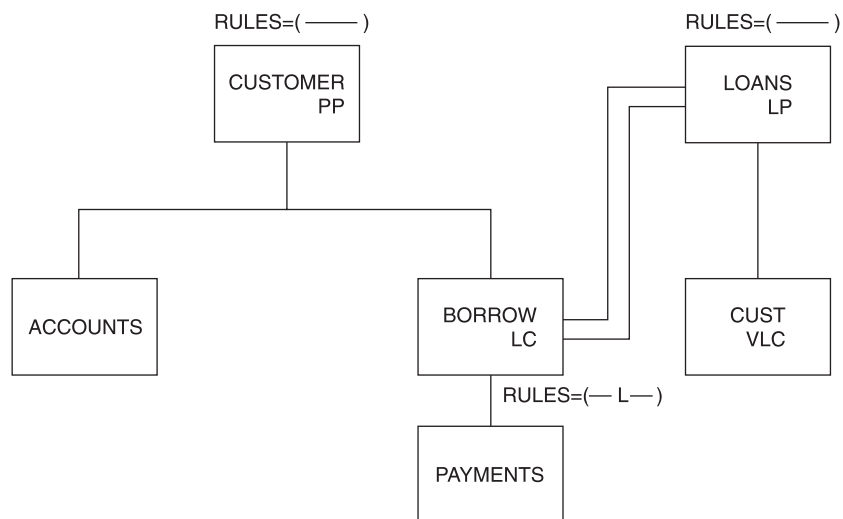


Figure 285. Logical Child, Virtual Pairing—Logical Delete Rule Example

```

GHU 'CUSTOMER  STATUS=' '
    'BORROW/LOANS'
DLET          STATUS=' '

GHU 'LOANS'      STATUS=' '
    'CUST/CUSTOMER'
DLET          STATUS=' '


```

Figure 286. Logical Child, Virtual Pairing—Logical Delete Rule Example: Calls and Status

The logical delete rule allows the logical child to be deleted physically or logically first. Physical dependents of the logical child are physically deleted, but they remain accessible from the logical path that is not logically deleted.

The delete of the virtual logical child sets the LD bit on in the physical logical child BORROW (BORROW is logically deleted).

Before



After




Figure 287. Logical Child, Virtual Pairing—Logical Delete Rule Example: Before and After

Delete Rules

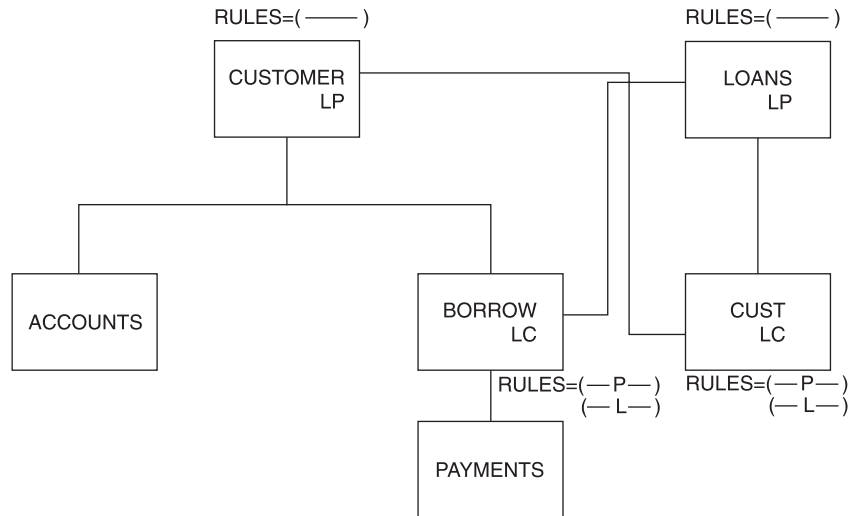


Figure 288. Logical Child, Physical Pairing—Physical or Logical Delete Rule Example

```

GHU 'CUSTOMER' STATUS=' '
    'BORROW/LOANS'
DLET STATUS=' '

GHU 'LOANS' STATUS=' '
    'CUST/CUSTOMER'
DLET STATUS=' '
  
```

Figure 289. Logical Child, Physical Pairing—Physical or Logical Delete Rule Example: Calls and Status

With the physical or logical delete rule, each logical child must be deleted from its physical path. Physical dependents of the logical child are physically deleted, but they remain accessible from the paired logical child that is not deleted.

Physically deleting BORROW sets the LD bit on in CUST. Physically deleting CUST sets the LC bit on in the BORROW segment.




Figure 290. Logical Child, Physical Pairing—Physical or Logical Delete Rule Example: Before and After

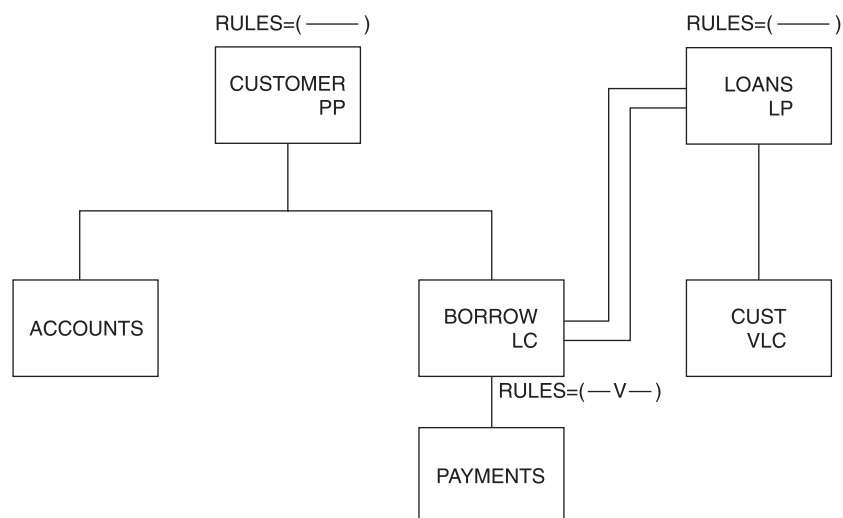


Figure 291. Logical Child, Virtual Pairing—Virtual Delete Rule Example

Delete Rules


```
GHU 'CUSTOMER  STATUS=' ' '  
    'BORROW/LOANS'  
DLET          STATUS=' ' '  
  
GHU 'LOANS'      STATUS='GE'  
    'CUST/CUSTOMER'
```

Figure 292. Logical Child, Virtual Pairing—Virtual Delete Rule Example: Calls and Status

The virtual delete rule allows the logical child to be deleted physically and logically. Deleting either path deletes both parts. Physical dependents of the logical child are physically deleted.

The previous delete deleted both paths because the delete rule is virtual. Deleting either path deletes both.

Before



After




Figure 293. Logical Child, Virtual Pairing—Virtual Delete Rule Example: Before and After

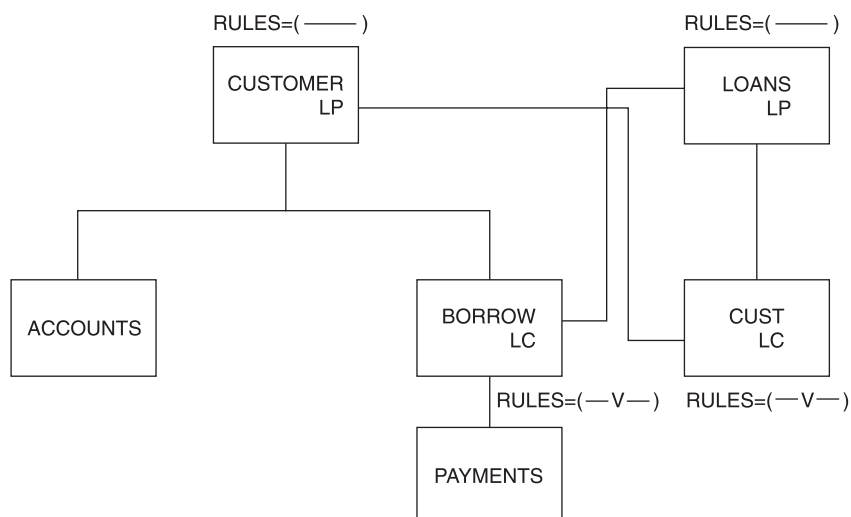


Figure 294. Logical Child, Physical Pairing—Virtual Delete Rule Example

```

GHU 'CUSTOMER  STATUS=' '
DLET      STATUS=' '

GHU 'LOANS'      STATUS='GE'
    'CUST/CUSTOMER'

```

Figure 295. Logical Child, Physical Pairing—Virtual Delete Rule Example: Calls and Status

With the virtual delete rule, deleting either logical child deletes both paired logical children. (Notice the PD and LD bit is set on in both.) Physical dependents of the logical child are physically deleted.

Physical dependents of the logical child are physically deleted.

Delete Rules




Figure 296. Logical Child, Physical Pairing—Virtual Delete Rule Example: Before and After

Accessibility of Deleted Segments

A physically deleted segment remains accessible under the following circumstances:

- A physical dependent of the deleted segment is a logical parent accessible from its logical children.
- A physical dependent of the deleted segment is a logical child accessible from its logical parent.
- A physical parent of the deleted segment is a logical child accessible from its logical parent. The deleted segment in this case is variable intersection data in a bidirectional logical relationship.

A logically deleted logical child cannot be accessed from its logical parent.

Neither physical or logical deletion prevents access to a segment from its physical or logical children. Because logical relationships provide for inversion of the physical structure, a segment can be physically or logically deleted or both, and still be accessible from a dependent segment because of an active logical relationship. A physically deleted root segment can be accessed when it is defined as a dependent segment in a logical DBD. The logical DBD defines the inversion of the physical DBD. Figure 297 shows the accessibility of deleted segments.'

When the physical dependent of a deleted segment is a logical parent with logical children that are not physically deleted, the logical parent and its physical parents are accessible from those logical children.

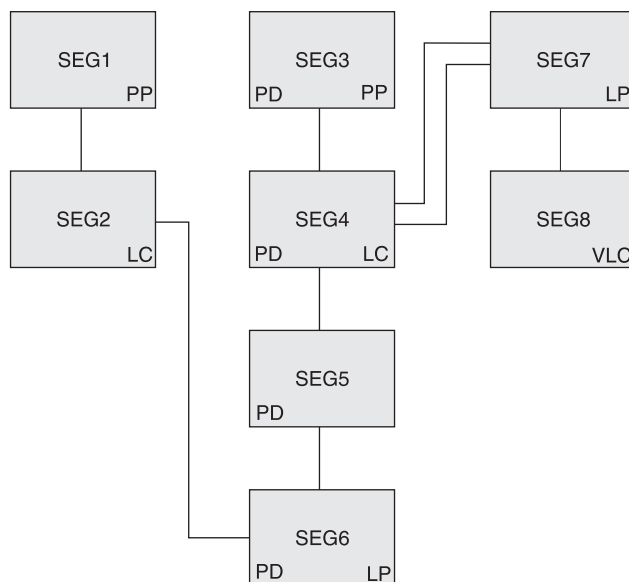


Figure 297. (Part 1 of 5). Example of Deleted Segments Accessibility

The physical structure in Figure 297 shows that SEG3, SEG4, SEG5, and SEG6 have been physically deleted, probably by issuing a DLET call for SEG3. This resulted in all of SEG3's dependents being physically deleted. (SEG6's delete rule is not P, or a 'DX' status code would be issued.)

SEG3, SEG4, SEG5, and SEG6 remain accessible from SEG2, the logical child of SEG6. This is because SEG2 is not physically deleted. However, physical dependents of SEG6 cannot be accessible, and their DASD space is released unless an active logical relationship prohibits

When the physical dependent of a deleted segment is a logical child whose logical parent is not physically deleted, the logical child, its physical parents, and its physical dependents are accessible from the logical parent.

The logical child segment SEG4 remains accessible from its logical parent SEG7 (SEG7 is not physically deleted). Also accessible are SEG5 and SEG6, which are variable intersection data. The physical parent of the logical child (SEG3) is also accessible from the logical child (SEG4).

A physically and logically deleted logical child can be accessed from its physical dependents (Figure 298 on page 502).

Delete Rules




Figure 298. (Part 2 of 5). Example of Deleted Segments Accessibility

The physical structure Figure 298 shows that logical child SEG4 is both physically and logically deleted.

From a previous example (part 1 of 4), we know SEG6 (a logical parent) is accessible from SEG2, if that segment (its logical child) is not physically deleted. We also know that once we've accessed SEG6, its physical parents (SEG5, SEG4, SEG3) are accessible. It doesn't matter that the logical child is logically deleted (which is the only difference between this example and that of part 1 of 4).

The third path cannot be blocked because no delete bit exists for this path. Therefore, the logical child SEG4 is accessible from its dependents even though it is been physically and logically deleted.

When a segment accessed by its third path is deleted, it is physically deleted in its physical data base, but it remains accessible from its third path (Figure 299 and Figure 300 on page 503).




Figure 299. (Part 3 of 5). Example of Deleted Segments Accessibility

```

GHU 'SEG5' STATUS=' '
DLET      STATUS=' '

```

Figure 300. (Part 4 of 5). Example of Deleted Segments Accessibility: Database Calls

SEG5 is physically deleted by the DLET call, and SEG 6 is physically deleted by propagation. SEG 2/SEG6 has unidirectional pointers, so SEG2 was considered logically deleted before the DLET call was issued. The LD bit is only assumed to be set on (Figure 301).

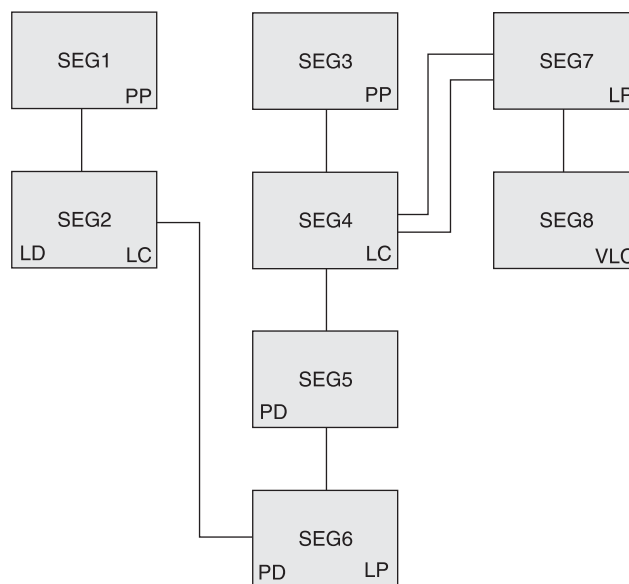


Figure 301. (Part 5 of 5). Example of Deleted Segments Accessibility

The results are interesting. SEG5 is inaccessible from its physical parent path (from SEG4) unless SEG4 is accessed by its logical parent SEG7 (SEG5 and SEG6 are accessible as variable intersection data). SEG5 is still accessible from its third path (from SEG6) because SEG6 is still accessible from its logical child. Thus, a segment can be physically deleted by an application program and still be accessible to that application program, using the same PCB used to delete the segment.

Possibility of Abnormal Termination

If a logical parent is physically and logically deleted, its DASD space is released. For this to occur, all of its logical children must be physically and logically deleted. However, the DASD space for these logical children cannot be released if the logical children have physical dependents with active logical relationships. Accessing such a logical child from its physical dependents (both the logical child and logical parent have been physically and logically deleted) can result in a user 850 through 859 abnormal termination if one of the following occurs:

- The LPCK is not stored in the logical child
- The concatenation definition is data sensitive to the logical parent

Figure 302 shows an example of abnormal termination.

Delete Rules

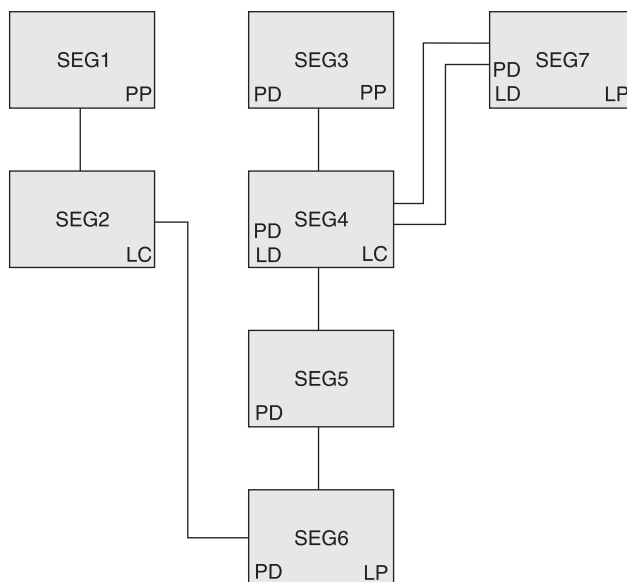


Figure 302. Example of Abnormal Termination

The logical parent SEG7 has been physically and logically deleted (the LD bit is never really set, but is assumed to be set. It is shown only for the purpose of illustration.) All of the logical children of the logical parent have also been physically and logically deleted. However, the logical parent has had its segment space released, whereas the logical child (SEG4) still exists. The logical child still exists because it has a physical dependent that has an active logical relationship that precludes releasing its space.

If an application program accesses SEG4 from its dependents (SEG1 to SEG2/SEG6 to SEG5), IMS must build the logical parent's concatenated key if that key is not stored in the logical child. When IMS attempts to access logical parent SEG7, abnormal termination will occur. The 850 through 859 abnormal termination codes are issued when a pointer is followed that doesn't lead to the expected segment.

Avoiding Abnormal Termination

You must avoid creating a physically deleted logical child that can be accessed from below in the physical structure (using its third path). A logical child can be accessed from below if any of its physical dependents are accessible through logical paths. Two methods exist in avoiding this situation.

- Method 1

The first method requires that logical paths to dependents be broken before the logical child is physically deleted. Breaking the logical path with method 1 is done using a P rule for the dependents as long as no physical deletes are propagated into the database. Therefore, no V rules on logical children can be allowed at or above the logical child, because, with the V rule, a propagated logical delete causes a physical delete without a P rule violation check. (For more information on this, see "Detecting Physical Delete Rule Violations" on page 505.) The L rule also causes propagation, if the PD bit is already set on, but the dependent's P rule will prevent that case. Similarly, no V rule can be allowed on any logical parent above the logical child, because the logical delete condition would cause the physical delete.

- Method 2

The second method requires breaking the logical path whenever the logical child is physically deleted. Breaking the logical path with this method is done for subordinate logical child segments using the V delete rule. Subordinate logical parent segments need to have bidirectional logical children with the V rule (must be able to reach the logical children) or physically paired logical children with the V rule. This method will not work with subordinate logical parents pointed to by unidirectional logical children.

Detecting Physical Delete Rule Violations

When a DLET call is issued, the delete routine scans the physical structure containing the segment to be deleted. The delete routine scans the physical structure to determine if any segment in it uses the physical delete rule and whether that rule is being violated. Figure 303 and Figure 304 show an example of violating the physical delete rule.

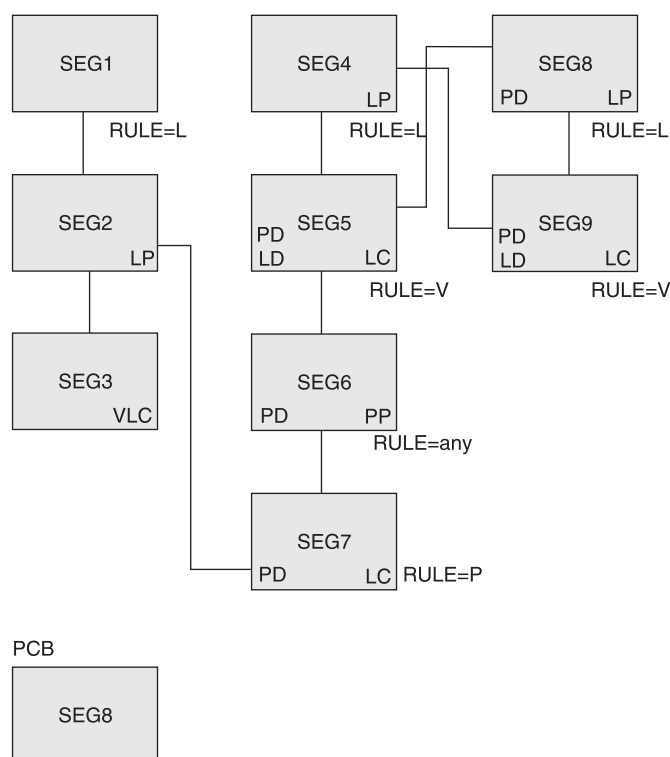


Figure 303. Example of Violation of the Physical Delete Rule

```

GHU 'SEG4' STATUS=' '
DLET      STATUS='DX'
  
```

Figure 304. Example of Violation of the Physical Delete Rule: Database Calls

SEG7 (the logical child of SEG2) uses the physical delete rule and has not been logically deleted (the LD bit has not been set on). Therefore, the physical delete rule is violated. A 'DX' status code is returned to the application program, and no segments are deleted.

Treating the Physical Delete Rule as Logical

If the delete routine determines that neither the segment specified in the DLET call nor any physical dependent of that segment in the physical structure uses the physical delete rule, any physical rule encountered later (logical deletion propagated to logical child or logical parent causing physical deletion—V rule—in another

Delete Rules

database) is treated as a logical delete rule. Figure 305 and Figure 306 show an example of treating the physical delete rule as logical.

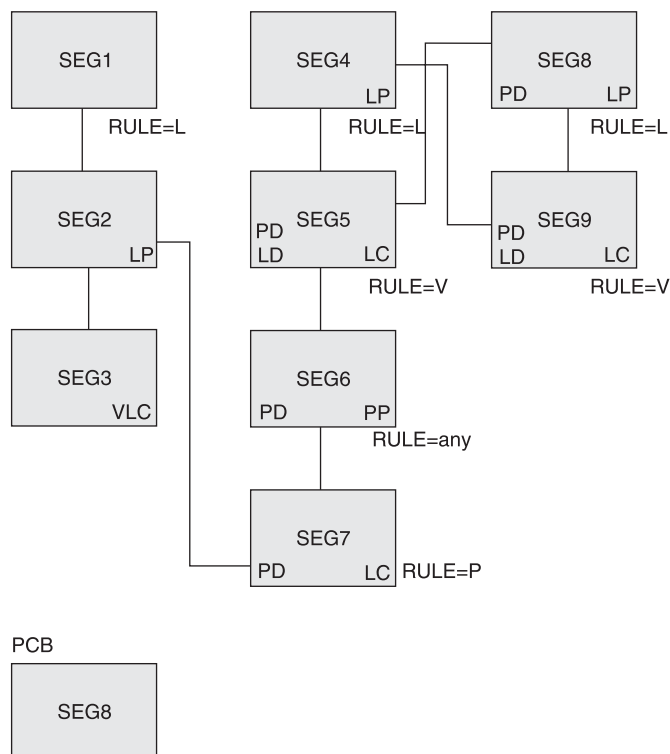


Figure 305. Example of Treating the Physical Delete Rule as Logical

```
GHU 'SEG8' STATUS=' '
DLET      STATUS=' '

```

Figure 306. Example of Treating the Physical Delete Rule as Logical: Database Calls

SEG8 and SEG9 are both physically deleted, and SEG9 is logically deleted (V rule). SEG5 is physically and logically deleted because it is the physical pair to SEG9 (with physical pairing setting the LD bit in one set, the PID bit in the other, and vice versa). Physically deleting SEG5 causes propagation of the physical delete to SEG5's physical dependents; therefore, SEG6 and SEG7 are physically deleted.

Note that the physical deletion of SEG7 is prevented if the physical deletion started by issuing a DLET call for SEG4. But the physical rule of SEG7 is treated as logical in this case.

Inserting Physically and Logically Deleted Segments

When a segment is inserted, a replace operation is performed (space is reused), and existing dependents of the inserted segment remain if:

- The segment to be inserted already exists (same segment type and same key field value for both the physical and logical sequencing)
- The delete bit is set on for that segment along the path of insertion

For HDAM and HIDAM databases, the logical twin chain is established as required, and existing dependents of the inserted segment remain.

For HISAM databases, if the root segment is physically and logically deleted before the insert is done, then the first logical record for that root in primary and secondary data set groups is reused. Remaining logical records on any OSAM chain are dropped.

Delete Rules Summary

The DLET Call

A DLET call issued against a concatenated segment (SOURCE=DATA/DATA, DATA/KEY, KEY/DATA) is a DLET call against the logical child only.

A DLET call against a logical child that has been accessed from its logical parent is a request that the logical child be logically deleted.

In all other cases, a DLET call issued against a segment is a request for that segment to be physically deleted.

Physical Deletion

A physically deleted segment cannot be accessed from its physical path, however, one exception exists: If one of the physical parents of the physically deleted segment is a logical child that can be accessed from its logical parent, then the physically deleted segment is accessible from that logical child. The physically deleted segments is accessible because the physical dependents of the logical child are variable intersection data.

Logical Deletion

By definition, a logically deleted logical child cannot be accessed from its logical parent. Unidirectional logical child segments are assumed to be logically deleted.

By definition, a logical parent is considered logically deleted when all its logical children are physically deleted and all its physical children that are part of a physically paired set are physically deleted.

Access Paths

Neither physical nor logical deletion of a segment prevents access to the segment from its physical or logical children, or from the segment to its physical or logical parents. A physically deleted root segment can be accessed only from its physical or logical children.

Propagation of Delete

In bidirectional physical pairing, physical deletion of one of the pair of logical children causes logical deletion of its paired segment. Likewise, logical deletion of one causes physical deletion of the other.

Physical deletion of a segment propagates logical deletion requests to its bidirectional logical children. Physical deletion of a segment propagates physical deletion requests to its physical children and to any index pointer segments for which it is the source segment.

Delete Rules

Further delete operations are governed by the following delete rules:

Logical Parent

When RULES=P is specified, if the segment is not already logically deleted, a DLET call against the segment or any of its physical parents results in a DX status code. No segments are deleted. If a request is made against the segment as a result of propagation across a logical relationship, then the P rule works like the L rule.

When RULES=L is specified, either physical or logical deletion can occur first, and neither causes the other to occur.

Delete Rules

When RULES=V is specified, either physical or logical deletion can occur first. If the segment is logically deleted as the result of a DLET call, then it is physically deleted also.

Physical Parent of a Virtually Paired Logical Child

RULES=P, L, or V is meaningless.

When RULES=B is specified and all physical children that are virtually paired logical children are logically deleted, the physical parent segment is physically deleted.

Logical Child

When RULES=P is specified, if the segment is not already logically deleted, then a DLET call requesting physical deletion of the segment or any of its physical parents results in a DX status code. No segments are deleted. If a delete request is made against the segment as a result of propagation across a logical relationship or if the segment is one of a physically paired set, then the rule works like the L rule.

When RULES=L is specified, either physical or logical deletion can occur first, and neither causes the other to occur.

When RULES=V is specified, either physical or logical deletion can occur first and either causes the other to occur. If this rule is used on only one segment of a physically paired set, it works like the L rule.

Space Release

Depending on the database organization, DASD space can or cannot be reused when it is released. DASD space for a segment is released when the following conditions are met:

- Space has been released for all physical dependents of the segment.
- The segment is physically deleted.
- If the segment is a logical child or a logical parent, then it is physically and logically deleted.
- If the segment is a dependent of a logical child (variable intersection data) and the DLET call was issued against a physical parent of the logical child, then the logical child is both physically and logically deleted.
- If the segment is a primary index pointer segment, the space is released for its target segment.

Insert, Delete, and Replace Rules Summary

Figure 307 summarizes rules by stating a desired result and then indicating the rule that can be used to obtain that result. Table 35 on page 509 lists the rules and how to specify them.




Figure 307. Insert, Delete, and Replace Rules Summary

Table 35. Specifying Insert, Delete, and Replace Rules

| Rule | RULES= Specification |
|-----------------------------------|----------------------|
| physical insert rule | RULES= (P,_,_) |
| logical insert rule | RULES= (L,_,_) |
| virtual insert rule | RULES= (V,_,_) |
| physical delete rule | RULES= (_,P,_) |
| logical delete rule | RULES= (_,L,_) |
| bidirectional virtual delete rule | RULES= (_,B,_) |
| virtual delete rule | RULES= (_,V,_) |
| physical replace rule | RULES= (_,_,P) |
| logical replace rule | RULES= (_,_,L) |
| virtual replace rule | RULES= (_,_,V) |

Insert Rules for Physical Parent Segment A: The insert rules for physical parent (PP) segment A control the insert of PP A using the logical path to PP A. The rules are as follows:

- To disallow the insert of PP A on its logical path, use the physical insert rule.
- To allow the insert of PP A on its logical path (concatenated with virtual logical child segment A), use either the logical or virtual rule.

Where PP A is already present, a logical connection is established to the existing PP A segment. The existing PP A can either be replaced or remain unchanged:

- If PP A is to remain unchanged by the insert call, use the logical insert rule.
- If PP A is to be replaced by the insert call, use the virtual insert rule.

Delete Rules for Physical Parent Segment A: The delete rules for PP segment A control the deletion of PP A using the logical path to PP A. The rules are as follows:

- To cause PP segment A to be deleted automatically when the last logical connection (through real logical child segment B to PP segment A) is broken, use the bidirectional virtual delete rule.
- The other delete rules for PP A are not meaningful.

Delete Rules

Replace Rules for Physical Parent Segment A: The replace rules for PP segment A control the replacement of PP A using the logical path to PP A. The rules are as follows:

- To disallow the replacement of PP A on its logical path and receive an 'RX' status code if the rule is violated by an attempt to replace PP A, use the physical replace rule.
- To disregard the replacement of PP A on its logical path, use the logical replace rule.
- To allow the replacement of PP A on its logical path, use the virtual replace rule.

Insert Rules for Logical Parent Segment B:

Note: These rules are identical to the insert rules for PP segment A.

The insert rules for logical parent (LP) segment B control the insert of LP B using the logical path to LP B. The rules are as follows:

- To disallow the insert of LP B on its logical path, use the physical insert rule.
- To allow the insert of LP B on its logical path (concatenated with virtual segment RLC B) use either the logical or virtual rule.

Where LP B is already present, a logical connection is established to the existing LP B segment. The existing LP B can either be replaced or remain unchanged:

- If LP B is to remain unchanged by the insert call, use the logical insert rule.
- If LP B is to be replaced by the insert call, use the virtual insert rule.

Delete Rules for Logical Parent Segment B: The delete rules for segment LP B control the deletion of LP B on its physical path. A delete call for a concatenated segment is interpreted as a delete of the logical child only. The rules are as follows:

- To ensure that LP B remains accessible until the last logical relationship path to that occurrence has been deleted, choose the physical delete rule. If an attempt to delete LP B is made while there are occurrences of real logical child (RLC) B pointing to LP B, a 'DX' status code is returned and no segment is deleted.
- To allow segment LP B to be deleted on its physical path, choose the logical delete rule. When LP B is deleted, it is no longer accessible on its physical path. It is still possible to access LP B from PP A through RLC B as long as RLC B exists.
- Use the virtual delete rule to physically delete LP B when it has been explicitly deleted by a delete call or implicitly deleted when all RLC Bs pointing to it have been physically deleted.

Replace Rules for Logical Parent Segment B:

Note: These rules are identical to the replace rules for PP segment A.

The replace rules for LP segment B control the replacement of LP B using the logical path to LP B. The rules are as follows:

- Use the physical replace rule to disallow the replacement of LP B on its logical path and receive an 'RX' status code if the rule is violated by an attempt to replace LP B.
- Use the logical replace rule to disregard the replacement of LP B on its logical path.
- Use the virtual replace rule to allow the replacement of LP B on its logical path.

Insert Rules for Real Logical Child Segment B: The insert rules do not apply to a logical child.

Delete Rules for Real Logical Child Segment B: The delete rules for RLC segment B apply to delete calls using its logical or physical path. The rules are as follows:

- Use the physical delete rule to control the sequence in which RLC B is deleted on its logical and physical paths. The physical delete rule requires that it be logically deleted before it is physically deleted. A violation results in a 'DX' status code.
- Use the logical delete rule to allow either physical or logical deletes to be first.
- Use the virtual delete rule to use a single delete call from either the logical or physical path to both logically and physically delete RLC B.

Replace Rules for Real Logical Child Segment B:

Note: These rules are identical to the replace rules for PP segment A.

The replace rules for LP B control the replacement of RLC B using the logical path to RLC B. The rules are as follows:

- Use the physical replace rule to disallow the replacement of RLC B on its logical path and receive an 'RX' status code if the rule is violated by an attempt to replace RLC B.
- To disregard an attempt to replace RLC B on its logical path, use the logical replace rule.
- To allow the replacement of RLC B on its logical path, use the virtual replace rule.

Delete Rules

Appendix C. Using OSAM as the Access Method

This appendix contains product-sensitive programming interface information.

You need to know the following information about OSAM if your database is using OSAM as an access method:

- OSAM is a special access method supplied with IMS.
- IMS communicates with OSAM using OPEN, CLOSE, READ, and WRITE macros.
- OSAM communicates with the I/O supervisor using the I/O driver interface.
- An OSAM data set can be read using either the BSAM or QSAM access method.
- The number of extents in an OSAM data set is limited by:
 - The maximum length of the data extent block (DEB)
 - The length of the sector number table that is created for rotational position sensing (RPS) devices

The length of a DEB is represented in a single byte that is expressed as the number of double words. The sector number table exists only for RPS devices and consists of a fixed area of eight bytes plus one byte for each block on a track, rounded up to an even multiple of eight bytes. A minimum-sized sector table (7 blocks per track) requires two double words. A maximum-sized sector table (255 blocks per track) requires 33 double words.

In addition, for each extent area (two double words), OSAM requires a similar area that contains device geometry data. Each extent requires a total of four double words. The format and length (expressed in double words) of an OSAM DEB are shown in Table 36.

Table 36. Length and Format of an OSAM DEB

| Format | Length |
|---------------------------------|------------------|
| Appendage sector table | 5 |
| Basic DEB | 4 |
| Access method dependent section | 2 |
| Subroutine name section | 1 |
| Standard DEB extents | 120 (60 extents) |
| OSAM extent data | 120 |
| Minimum sector table | 2 |

With a minimum-sized sector table, the DEB can reflect a maximum of 60 DASD extents. With a maximum-sized sector table, the DEB can reflect a maximum of 52 DASD extents.

- An OSAM data set can be opened for update in place and extension to the end through one data control block (DCB). The phrase “extension to the end” means that records can be added to the end of the data set and that new direct-access extents can be obtained.
- An OSAM data set does not need to be formatted before use.
- An OSAM data set can use fixed-length blocked or unblocked records.
- The maximum size of an OSAM data set depends on the block size of the data set and whether it is a HALDB OSAM data set. The size limits for OSAM data sets are:
 - 8 GB for a non-HALDB OSAM data set that has an even-length block size

OSAM as the Access Method

- 4 GB for a non-HALDB OSAM data set that has an odd-length block size
- 4 GB for a HALDB OSAM data set
- File mark definition is always used to define the current end of the data set. When new blocks are added to the end of the data set, they replace dummy pre-formatted (by OSAM) blocks that exist on a logical cylinder basis. A file mark is written at the beginning of the next cylinder, if one exists, during a format logical cylinder operation. This technique is used as a reliability aid while the OSAM data set is open.
- OSAM EXCP counts are accumulated during OSAM End of Volume (EOV) and close processing.
- Migrating OSAM data sets utilizing ADRDSSU and the DFSMSdss™ component of z/OS DFSMS: DFSMSdss will migrate the tracks of a data set up to the last block written value (DS1LSTAR) as specified by the DSCB for the volume being migrated. If the OSAM data set spans multiple volumes that have not been pre-allocated, the DS1LSTAR field for each DSCB will be valid and DFSMSdss can correctly migrate the data.

If the OSAM data set spans multiple volumes that have been pre-allocated, the DS1LSTAR field in the DSCB for each volume (except the last) can be zero. This condition will occur during the loading operation of a pre-allocated, multi-volume data set. The use of pre-allocated volumes precludes EOV processing when moving from one volume to another, thereby allowing the DSCBs for these volumes not to be updated. The DSCB for the last volume loaded is updated during close processing of the data set.

DFSMSdss physical DUMP or RESTORE commands with the parameters ALLEXCP or ALLDATA must be used when migrating OSAM data sets that span multiple pre-allocated volumes. These parameters will allow DFSMSdss to correctly migrate OSAM data sets.

Related Reading: For more information on the z/OS DFSMSdss component of DFSMS and the ALLEXCP and ALLDATA parameters of the DUMP and RESTORE commands, see the *DFSMSdss Storage Administration Reference*.

Other z/OS access methods (VSAM and SAM) are used in addition to OSAM for physical storage of data.

Related Reading: For information about defining OSAM subpools, see *IMS Version 9: Installation Volume 2: System Definition and Tailoring*.

Appendix D. Correcting Bad Pointers

Ordinarily, bad pointers should not occur in your database. When they do, the cause is typically:

- Failure to run database backout
- Failure to perform emergency restart
- Omitting a log during backout or recovery

The normal way to correct a bad pointer is to perform recovery. However, some cases exist in which a bad pointer can be corrected through reorganization. A description of the circumstances in which this can or cannot be done is as follows:

- PC/PT pointers. The HD Unload utility issues unqualified GN calls to read a database. If the bad pointer is a PC or PT pointer, DL/I will follow the bad pointer and the GN call will fail. Therefore, reorganization cannot be used to correct PC or PT pointers.
- LP/LT pointers. LP and LT pointers are rebuilt during reorganization. However, DL/I can follow the LP pointer during unload. If the logical child segment contains a direct LP pointer and the logical parent's concatenated key is not physically stored in the logical child segment, DL/I follows the bad LP pointer to construct the logical parent's concatenated key. This causes an ABEND.
- LP pointer. When DBR= is specified for pre-reorganization and the database has direct LP pointers, the HD Unload utility saves the old LP pointer. Bad LP pointers produce an error message (DFS879) saying a logical child that has no logical parent exists.
- LP pointer. When DBIL= is specified for pre-reorganization of a logical child or parent database, the utilities that resolve LP pointers use concatenated keys to match logical parent and logical child segments. New LP pointers are created.

Appendix E. HALDB Partition Definition utility

The HALDB Partition Definition utility is an ISPF application that allows you to manage the partitions of an IMS HALDB. This utility can be used in place of the DBRC INIT.DB and INIT.PART commands to register the HALDB master in the RECON data set. The HALDB master is registered at the time its first partition is defined.

Related Reading: For more information on how HALDBs are maintained in the RECON, see *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

Important: The HALDB Partition Definition utility will not impact online IMS subsystems with regard to RECON contention. The RECON is only reserved for the time it takes to process a DBRC request. It is not held for the duration of the utility execution.

To access the HALDB utility:

1. Log on to TSO.
2. Start ISPF.
3. From the ISPF command line, type: `tso %dfshaldb` and press Enter.

The utility consists of several panels and programs that perform various actions on the HALDB and its partitions.

Important: The Panel IDs are shown enclosed in parentheses in the caption of each panel image here. To enable Panel IDs to be displayed in the upper left corner of each of your panels; enter `panelid` on the ISPF command line and press Enter.

In this Appendix:

- “The Partitioned Databases Panel” on page 518
- “Accessing Help Information” on page 519
- “Exiting the Utility” on page 519
- “Displaying the ISPF Member List” on page 520
- “Opening HALDB Partitions” on page 521
- “Defining Data Set Group Information” on page 533
- “Displaying the List of Defined Partitions” on page 534
- “Opening Database Information” on page 542
- “Deleting Database Information” on page 543
- “Exporting Database Information” on page 543
- “Importing Database Information” on page 544
- “Displaying the IMS Concatenation” on page 544
- “Selecting an IMS Configuration Using the Partition Definition Utility” on page 545
- “Using Batch to Export or Import Partition Information” on page 547
- “DSPXRUN Command Syntax” on page 548

The Partitioned Databases Panel

You define the HALDB that you want to manipulate on the Partitioned Databases panel. Here you specify the type of action to perform, for example: define, modify, or view. The succeeding panels guide you through the processes.

The Partitioned Databases panel has *point-and-shoot* text fields (in turquoise by default). To use the point-and-shoot fields, just position the cursor on the text and press the enter key.

The Figure 308 on page 518 provides space for you to enter a HALDB name, allowing HALDB to gather information about that HALDB. The information can be retrieved from DBDLIB or from RECON depending on the option you select and the current state of the partitions. Following Figure 308 on page 518 are descriptions of the panel fields.

Help

Partitioned Databases

Type a database name and choose an option. Then press Enter.
To select a database from a list, type a filter (*) and press F4.

Configuration . . . : DEFAULT

Database name . . . IVPDB1 +

Option —

1. OPEN DATABASE partitions
2. Open database information
3. Delete database information
4. Export database definitions
5. Import database definitions
6. Show IMS DDname concatenation
7. Select IMS RECON / DBDLIB libraries

To exit the application, press F3.

Command ==>
F1=Help F3=Exit F4=Prompt

Figure 308. Partitioned Databases panel (DSPXPAA)

The options in Figure 308 allow you to perform the following actions:

1. Create or change HALDB partitions (see “Opening HALDB Partitions” on page 521 and “Displaying the List of Defined Partitions” on page 534).
2. View or change HALDB information (see “Opening Database Information” on page 542).
3. Delete HALDB information (see “Deleting Database Information” on page 543).
4. Export HALDB information (see “Exporting Database Information” on page 543).
5. Import HALDB information (see “Importing Database Information” on page 544).
6. Show the IMS concatenation (see “Displaying the IMS Concatenation” on page 544).
7. Select an IMS configuration (see “Selecting an IMS Configuration Using the Partition Definition Utility” on page 545).

Configuration

The configuration is a name you have specified that identifies a set of DBD libraries and a set of RECON data sets. If you already have the IMS DD statement allocated from the logon procedure and if you have the

IMS.SDFSRESLs allocated to the STEPLIB DD statement, you do not need to use the **Configuration** option. If you do define and select a configuration, those data sets will override the allocations from the logon procedure.

Database name

Enter up to 8 alphanumeric characters (the first character must be alphabetic). The HALDB name must be a member from a DBDLIB data set. DBDLIB data sets must be allocated under a DD name of IMS. The database name that you specify is remembered across ISPF sessions.

You can include an asterisk to indicate that you want a member list display. The asterisk can appear alone or as part of the name to limit the list that is displayed. (see “Displaying the ISPF Member List” on page 520)

Important: When you include an asterisk as part of the member name, the concatenation for the 'IMS' DD name may contain only up to 4 data sets. This is an ISPF restriction.

Option

A numeric value that indicates the type of processing to perform. The number corresponds to one of the actions in the list.

Accessing Help Information

The Partitioned Databases panel has help information available from the action bar (Figure 309). This information is available from other panels as well.

Help

1. Help for Help
2. Extended help
3. Keys help
4. Help Index
5. About

Figure 309. Help Action Bar Choices

Help information can also be obtained by pressing the help key. The help displayed depends on the circumstances and on the placement of the cursor when the help key is pressed.

- If an error message is displayed, more information on the error might be displayed.
- If the cursor is on an input field, information about the field is displayed, otherwise information about the panel is displayed.

Important: The F1 key is set to invoke the Help dialogs.

Exiting the Utility

The panels in the HALDB Partition Definition utility support an exit key (F3) which will return your session to a display of the Figure 308 on page 518. When you use the exit key from Figure 308 your session will exit the HALDB Partition Definition utility altogether. Since you can accidentally press the exit key, a confirmation panel, shown in Figure 310 on page 520, will allow you to continue without losing unsaved changes.

Exiting the Utility

Exit Confirmation

—

1. Exit without saving.

2. Do not exit.

F1=Help F12=Cancel

Figure 310. Exit Confirmation Panel

To exit pull-down panels press the cancel key (F12), then the exit key (F3) if you wish to leave the HALDB Partition Definition utility panels altogether.

Displaying the ISPF Member List

When you include an asterisk in the database name field, a member list for the members of the IMS DD name concatenation with a name that matches the filter is displayed. A sample member list display is shown in Figure 311.

| File Help | | | | | | | |
|----------------------------|-----------------|--------|----------|----|----|----|---------------------|
| ----- | | | | | | | |
| MEMBER LIST | IMSIVP81.DBDLIB | | | | | | Row 00001 of 00011 |
| Name | Size | TTR | Alias-of | AC | AM | RM | ---- Attributes --- |
| . DBFSAMD1 | 00000158 | 00013B | | 00 | 24 | 24 | |
| . DBFSAMD2 | 000001A0 | 000143 | | 00 | 24 | 24 | |
| . DBFSAMD3 | 000006E0 | 00014B | | 00 | 24 | 24 | |
| . DBFSAMD4 | 000002C8 | 000207 | | 00 | 24 | 24 | |
| . DI21PART | 00000230 | 000133 | | 00 | 24 | 24 | |
| . IVPDB1 | 00000138 | 000103 | | 00 | 24 | 24 | |
| . IVPDB1I | 00000138 | 00010B | | 00 | 24 | 24 | |
| . IVPDB2 | 00000130 | 000113 | | 00 | 24 | 24 | |
| . IVPDB3 | 00000188 | 00011B | | 00 | 24 | 24 | |
| . IVPDB4 | 00000110 | 000123 | | 00 | 24 | 24 | |
| . IVPDB5 | 000000B0 | 00012B | | 00 | 24 | 24 | |
| **End** | | | | | | | |
| Command ==> | | | | | | | |
| Scroll ==> CSR | | | | | | | |
| F1=Help F3=Exit F12=Cancel | | | | | | | |

Figure 311. ISPF Member List Display (DSPXPAM)

The member list originates from the PDS directories of the IMS concatenation. The members that are displayed can be HALDB or non-HALDB. The member list is a standard ISPF list so there is no IMS-specific information displayed.

From the member list, you can select the HALDB name to process by typing in the far-left column. If the name selected is not for a partitioned database, an error message is displayed. You can select a HALDB name with the slash (/) character and the File action to select the type of actions to perform. The same actions that are shown on Figure 308 on page 518 are available here.

If you specify an option on the Partitioned Databases panel (518), you do not need to use the File Action bar; just press Enter. You can use the File Action bar to override the option that you specified an option on the Partitioned Databases panel.

The list of HALDBs in the Member List panel can be manipulated by using the File action bar (Figure 312 on page 521).

File Help

1. Open partition
2. Open database
3. Delete database
4. Export database
5. Import database

Figure 312. File Action Bar Choices

The options on the File Action bar allow you to perform the following actions:

- Create or change HALDB partitions (see “Opening HALDB Partitions” and “Displaying the List of Defined Partitions” on page 534).
- View or change HALDB information (see “Opening Database Information” on page 542).
- Delete HALDB information (see “Deleting Database Information” on page 543).
- Export HALDB information (see “Exporting Database Information” on page 543).
- Import HALDB information (see “Importing Database Information” on page 544).

Opening HALDB Partitions

Before you can define the partitions for a HALDB, you must use the DBDGEN process to define the HALDB as a partitioned database.

The first time you choose a HALDB you must set values for the HALDB master; see Figure 313 on page 522. When you press Enter to continue, you set the defaults for the partitions, see Figure 314 on page 524. When you press Enter to continue again, you define partitions using those defaults. You can modify each partition uniquely as they are created or you can modify them later from the list of partitions. Figure 316 on page 530 shows an example of the panel to specify the partition information.

After the initial set of partitions is defined (and whenever you select that HALDB again), you will see the Database Partitions display (see Figure 321 on page 535 in Displaying the List of Defined Partitions).

Important: Most of the information initially displayed on the panel in Figure 313 on page 522 is extracted from the DBDLIB member. You can change the displayed information, but that information is not saved back into the DBDLIB member (the definition is saved in the RECON data sets).

Each HALDB can support up to 1001 partitions.

```
Help
-----
                        Partitioned Database Information

Type the field values.  Then press Enter to continue.

Database name . . . . . : IVPDB1

                        Master Database values
Part. selection routine . . . DFSIVD1
RSR global service group . . . BKUPGRP1
RSR tracking type . . . . . DBTRACK
Share level . . . . . 0
Database organization . . . : PHDAM
Recoverable? . . . . . Yes
Number of data set groups . : 10
Online Reorganization Capable: Yes

To exit the application, press F3.

Command ==>
F1=Help    F3=Exit    F12=Cancel
```

Figure 313. Partitioned Database Information (DSPXPOA)

The following are descriptions of the fields on the Partitioned Database information screen:

Database name

Enter 1 to 8 alphanumeric characters. This is the name you selected from the previous panel (see Figure 308 on page 518); it is the name of the HALDB that you are defining.

Part. selection routine

Enter 1 to 8 alphanumeric characters (the first character must be alphabetic). This is the name of the Partition Selection Exit Routine provided by you.

RSR global service group

Enter 1 to 8 alphanumeric characters (the first character must be alphabetic). This is an optional parameter used to specify the RSR global service group that the HALDB is to be assigned to.

RSR tracking type

This is an optional parameter you use to specify the type of RSR tracking (shadowing) for a partition assigned to a global service group. The type, RCVTRACK or DBTRACK, cannot be specified without an RSR global service group having been defined for the HALDB master.

- DBTRACK- indicates HALDB readiness tracking is to be done.
- RCVTRACK- indicates recovery readiness tracking is to be done.

DBTRACK is the default.

Share level

0, 1, 2, or 3. Share level is an optional parameter you use to specify the level of data sharing that authorized subsystems can share a HALDB at.

Share level 0 is the default.

Database organization

This field indicates the type of HALDB organization, you can specify either: PSINDEX, PHIDAM, or PHDAM.

Recoverable?

Yes indicates that the HALDB is recoverable. No indicates that the HALDB is not recoverable. Yes is the default. If an RSR global service group is specified, the recoverable field must be Yes.

Related Reading: For more information on non-recoverable databases see the *IMS Version 9: Operations Guide*.

Number of data set groups

This is the number of data sets in the groups that contain data as specified in the DBDGEN.

Online Reorganization Capable

Yes specifies that this HALDB supports online reorganization. No specifies that this HALDB does not support online reorganization. These specifications are stored in the DBRC RECON data set.

Related Reading:

- For more information on reorganizing HALDBs online, see “HALDB Online Reorganization” on page 370.
- For more information on DBRC and the RECON data set, see *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

Figure 314 on page 524 shows the partition default information.

Opening HALDB Partitions

```
Help
-----
                        Partition Default Information

Type the field values.  Then press Enter to continue.

Database name . . . . . : IVPDB1

                        Processing options
Automatic definition . . . . No
Input dataset . . . . . : 'IMS.IVPDB1.KEYS'
Use defaults for DS groups. . No

                        Defaults for partitions
Partition name . . . . . IVPD101
Data set name prefix . . . . IMS.DB01.FINANCE.YEAR1998.CURR

Randomizer
Module name . . . . . DD41DUP2
Anchor . . . . . 2
High block number. . . . . 999
Bytes . . . . . 2000

Free Space
Free block freq. factor. . 0
Free space percentage. . . 0

                        Defaults for data set groups
Block Size . . . . . 8192

DBRC options
Max. image copies. . . . . 2
Recovery period. . . . . 0
Recovery utility JCL . . . RECOVJCL
Default JCL. . . . .
Image copy JCL . . . . . ICJCL
Online image copy JCL. . . OICJCL
Receive JCL. . . . . RECVJCL
Reusable? . . . . . No

To exit the application, press F3.

Command ==>
F1=Help  F3=Exit  F6=Groups  F12=Cancel
```

Figure 314. Partition Default Information (DSPXPCA)

Important:

- The Randomizer section is present only if the HALDB is PHDAM.
- The Defaults for data set groups section is present only if there is only one data set group specified during DBDGEN. If there are multiple data set groups, use F6=Groups to display all data set groups using the dialog described in “Defining Data Set Group Information” on page 533.

The following are descriptions of the fields on the Partition Default Information screen:

Database name

This is the name you selected from the previous panel (see Figure 308 on page 518), it is the name of the HALDB that you are defining.

Automatic definition

The value can be Yes or No. Specifying yes will cause the partitions to be defined automatically based on your choices for partition name (that must

include percent sign characters for placeholders. see “Automatic Partition Definition” on page 527) and input data set.

Specifying No allows you to specify unique values for each partition.

Yes is the default.

Input data set

Provide the name of a z/OS data set. Specify a member name if it is a PDS. Each line of the data set must contain a partition selection string or the high key value to be used during partition definition.

Related Reading: See “Automatic Partition Definition” on page 527 and “Manual Partition Definition” on page 528 for more details on defining partitions.

Use defaults for DS groups

This value can be Yes or No. This option determines if all data set groups are automatically set to the same defaults or if you are prompted to provide values for each group. It can be left blank if *automatic definition* is set to Yes.

Partition name

Enter 1 to 7 alphanumeric characters (the first character must be alphabetic). The Partition name is used as a prefix to the DDNAMEs of its data sets, and so it must be unique.

Related Reading: For automatic definitions, you need to include percent signs (%) as placeholders for an alphanumeric sequence number (A-Z, 0-9). See “Automatic Partition Definition” on page 527 for more details.

Data set name prefix

Any alphanumeric name that is valid in JCL with a maximum length of 37 characters.

Module name

Enter 1 to 8 alphanumeric characters (the first character must be alphabetic). This is the name of the randomizing module. A randomizing module controls root segment placement in, or retrieval from, the PHDAM HALDB. This parameter is for PHDAM HALDBs only.

Anchor

1 to 3 numeric digits, with a range of 1 to 255. Specifies the number of root anchor points desired in each control interval or block in the root addressable area of a PHDAM HALDB. The anchor operand must be an unsigned decimal integer and must not exceed a value of 255. Typical values are from 1 to 5. This parameter is for PHDAM HALDBs only.

The default value of this parameter is 1.

High block number

A numeric unsigned decimal integer value with a range of 0 to $2^{24} - 1$. This value specifies the maximum relative block number value that the user wishes to allow a randomizing module to produce for this HALDB. This parameter is for PHDAM HALDBs only. This value determines the number of control intervals or blocks in the root addressable area of an PHDAM HALDB.

A high block number of zero means that no upper limit check is performed on the RBN created by the randomizing module. That is, it is all root addressable area.

Bytes A numeric unsigned decimal integer value with a range of 1 to $2^{24} - 1$.

Opening HALDB Partitions

This value specifies the maximum number of bytes of a HALDB record that can be stored into the root addressable area in a series of inserts unbroken by a call to another HALDB record.

A value of 0 (zero) means that all bytes are addressable. It is equivalent to omitting the bytes parameter from the RMNAME keyword in the DBD macro statement in DBDGEN. This parameter is for PHDAM HALDBs only.

Related Reading: For more information on the DBD macro statement in DBDGEN, see *IMS Version 9: Utilities Reference: System*.

Free block freq. factor

A numeric unsigned decimal integer from 0 to 100, except 1. The free block frequency factor (fbff) specifies that every *n*th control interval or block in this data set group is left as free space during HALDB load or reorganization (where fbff=*n*). The range of fbff includes all integer values from 0 to 100 except fbff=1. The default value for fbff is 0.

Free space percentage

Two numeric unsigned decimal integer digits with a range from 0 to 99. The *fspf* is the free space percentage factor. It specifies the minimum percentage of each control interval or block that is to be left as free space in this data set group.

The default value for *fspf* is 0.

Block size

A numeric unsigned even decimal integer with a range from 1 to 32,000. *The block size value is used by OSAM only.* An initial value of 4096 is displayed. If the HALDB is not OSAM, the block size field is not displayed.

Related Reading: For more information on the INIT.DBDS command, see *IMS Version 9: Database Recovery Control (DBRC) Guide and Reference*.

Max. image copies

A required parameter you use to specify the number of image copies that DBRC maintains for the identified DBDS. The value must be a unsigned decimal integer from 2 to 255.

Recovery period

An optional parameter you use to specify the recovery period of the image copies for the specified DBDS.

Specify an unsigned decimal integer from 0 to 999 that represents the number of days that information about the image copies is kept in RECON. If you specify 0, there is no recovery period. 0 is the default.

Recovery utility JCL

Enter 1 to 8 alphanumeric characters (the first character must be alphabetic). This is an optional parameter you use to specify the name of a member of a partitioned data set of skeletal JCL. When you issue the GENJCL.RECOV command, DBRC uses this member to generate the JCL to run the Database Recovery utility for the identified DBDS.

RECOVJCL is the default member name.

Default JCL

Enter 1 to 8 alphanumeric characters (the first character must be alphabetic). This is an optional parameter you use to specify an implicit skeletal JCL default member for the DBDS. The specified member is used by the GENJCL.IC, GENJCL.OIC, and GENJCL.RECOV commands to resolve keywords that you have defined.

Image copy JCL

Enter 1 to 8 alphanumeric characters (the first character must be alphabetic). This is an optional parameter you use to specify the name of a member of a partitioned data set that contains skeletal JCL. When you issue the GENJCL.IC command, DBRC uses this member to generate the JCL to run the Database Image Copy utility for the identified DBDS.

ICJCL is the default member name.

Online image copy JCL

Enter 1 to 8 alphanumeric characters (the first character must be alphabetic). This is an optional parameter you use to specify the name of a member of a partitioned data set that contains skeletal JCL. DBRC uses this member when you issue the GENJCL.OIC command to generate the JCL to run the Online Database Image Copy utility for the identified DBDS.

OICJCL is the default member name.

Receive JCL

Enter 1 to 8 alphanumeric characters (the first character must be alphabetic). This is an optional parameter you use to specify the name of the skeletal JCL member used by the GENJCL.RECEIVE command.

RECVJCL is the default member name.

Reusable?

The value is either Yes or No. Specifies whether the Database Image Copy utility, or the Online Database Image Copy utility are to reuse previously defined image copy data sets.

No is the default value.

Automatic Partition Definition

In the Partition Default Information panel (see Figure 314 on page 524) you can set *Automatic definition* to yes and have your partitions defined without intervention. You must have previously created a data set and it must contain your partition selection strings. Specify the name of the data set in the *input data set* field of the panel depicted in Figure 314 on page 524.

Each line of the input data set must contain a partition selection string or the high key value to be used during partition definition. The file must contain only one value on each line of the file, with the value left-justified. The length of the string is determined by the last non-blank character. Each record must contain only one string.

In the partition name field, include percent signs (%) as placeholders for an alphanumeric sequence number (A-Z, 0-9). If you type a partition name like:

Partition name IVPD1%%

The partitions are created in the following sequence:

```
IVPD1AA
IVPD1AB
IVPD1AC
.
.
IVPD1AZ
IVPD1A0
IVPD1A1
IVPD1A2
.
```

Opening HALDB Partitions

```
.  
IVPD1A9  
IVPD1BA  
IVPD1BB  
IVPD1BC  
.  
.
```

When you press Enter, as many partitions as you have key values in the input data set are automatically generated.

If you want to generate partition names that will allow you to preserve your naming sequence when expanding your database in the future, you can specify a partition name like IVP1%%A. The partitions would then be created in the following sequence:

```
IVP1AAA  
IVP1ABA  
.  
.  
IVP1AZA  
IVP1A0A  
IVP1A1A  
.  
.  
IVP1A9A  
IVP1BAA  
.  
.
```

When automatic definition is processing, a status panel is displayed (Figure 315 on page 528). This automatic definition status panel is updated as new partitions are defined.

| |
|--|
| Automatic definition status |
| Partitions defined: 3 of 166 |
| Elapsed time: 00:02 |

Figure 315. Automatic Definition Status

After automatic definition is complete, in the Database Partitions panel (Figure 321 on page 535) you can see that the partition selection string is filled-in with information from your input data set.

Manual Partition Definition

On the Partition Default Information panel (Figure 314 on page 524) you can set **Automatic definition** to No so that you can define the partitions serially. You can still use an input data set even though you set **Automatic definition** to No.

- If you specify an input data set, you must have previously created the data set and it must contain your partition selection strings. The partition selection string field (in Figure 316 on page 530) is primed from your input data set. For each partition, the partition selection string is filled-in from a record of the input data set. If you try to define more partitions than there are key values, the last key value from the input data set is displayed on the Change Partition panel (Figure 316 on page 530) and you will have to change it manually.
- If you do not specify an input data set to provide the partition high key values, the partition high key values can be added manually for each partition.

- If you did not specify a partition selection exit, the partition high key values are required.
- If you did specify a partition selection exit, the partition selection string values are optional.

After you set the defaults and press the enter key, the partition definition screen is displayed. You can modify the fields and press the enter key to define the partition. After you press the enter key, the partition is defined in RECON and the partition definition panel is displayed again so that you can define more partitions. The partition ID is incremented each time a partition is defined. Press the cancel key (PF12) to prevent the displayed partition from being defined.

When you press PF12 to stop defining new partitions, the Partitioned Databases panel (Figure 308 on page 518) is displayed again. You may also choose to stop defining new partitions by pressing F11=List; a list of defined partitions (see “Displaying the List of Defined Partitions” on page 534) is displayed.

Help

Change Partition

Type the field values. Then press Enter.

Database name IVPDB1
Partition name IVPD101
Partition ID. 1
Data set name prefix. . . . IMS.DB01.FINANCE.YEAR1998.CURR
Partition Status.

Partition Selection String

| | | | | | | |
|-----|----------|----------|----------|----------|------------------|-------------|
| +00 | F2F0F0F3 | 4BF2F2F4 | 40F1F77A | F2F57AF0 | 2003.224 | 17:25:0 |
| +10 | F94BF6F3 | F3F12432 | 00000000 | 00001020 | 9.6331 | |
| +20 | A840C1A5 | 85404040 | 40E28195 | 40D196A2 | y Ave | San Jos |
| +30 | 856B40C3 | C14040F9 | F5F1F4F1 | 00100020 | e, CA | 95141.... |
| +40 | 00050000 | 40F0F34B | F0F3F440 | 00000100 | | 03.034 |
| +50 | F1F8F0F0 | C9C2D4E2 | C5D9E540 | 40C9C2D4 | 1800IBMSERV | IBM |
| +60 | 40C39699 | 974B4040 | F5F5F540 | C2818993 | Corp. | 555 Bail |
| +70 | A840C1A5 | 85404040 | 40E28195 | 40D196A2 | y Ave | San Jos |
| +80 | 856B40C3 | C14040F9 | F5F1F4F1 | 00403010 | e, CA | 95141. .. |
| +90 | 00010500 | 40F0F34B | F2F4F340 | 00324020 | | 03.243 .. |
| +A0 | 9201913C | D2FE933D | 913C1F66 | 4360A005 | k.j.K.l.j.... | -.. |
| +B0 | 3233A200 | D996A281 | 6BD785A3 | 85996B40 | ..s.Rosa,Peter, | |
| +C0 | 000080D4 | 81A3A3F9 | 71C4C6F8 | F1F4C6C2 | ...Matt9.DF814FB | |
| +D0 | 9311913C | F6F4F8F6 | 943C1F66 | 4360A005 | l.j.6486m.... | -.. |
| +E0 | 41E3453C | 06000045 | 10110220 | 10416220 | .T..... | |
| +F0 | FFFFF900 | 00004920 | 18007410 | 94000300 | ..9..... | m... |

Randomizer
Module name DD41DUP2
Anchor 2
High block number. . . . 999
Bytes 2000

Free Space
Free block freq. factor. . 0
Free space percentage. . . 0

Attributes for data set group A
Block Size 8192

DBRC options
Max. image copies. . . . 2
Recovery period. 0
Recovery utility JCL . . RECOVJCL
Default JCL.
Image copy JCL ICJCL
Online image copy JCL. . OICJCL
Receive JCL. RECVJCL
Reusable? No

Command ==>
F1=Help F3=Exit F5=String F6=Groups F12=Cancel

Figure 316. Change Partition (DSPXPPA)

Important:

- The Randomizer section is present only if the HALDB is PHDAM.
- The data set group attributes section is present only if there is only one data set group specified during DBDGEN. If there is more than one data set group, use F6=Groups to display all data set groups using the dialog described in “Defining Data Set Group Information” on page 533.

The following are descriptions of the fields on the Change Partition screen:

Partition ID

A numeric value between 1 and 32 767, but less than the current high partition ID value for this HALDB. The Partition Definition utility generates the partition ID for you, regardless of whether you create your partitions manually or automatically. DBRC records this number in the RECON data set. Data set names include the partition ID of the partition to which they belong.

After an ID is assigned to a partition, you cannot change it.

Partition Status

You can disable a partition by typing `disable` in the **Partition Status** field. Usually, you would only disable a partition prior to deleting it.

To enable a disabled partition, type `enable` in the **Partition Status** field.

Partition High Key

The Partition High Key field allows you to specify the highest database record root key that a partition can contain. The partition high key is determined by your installation. IMS treats the partition high key as a hexadecimal value. You must enter a value in the Partition High Key field.

The length of the Partition High Key field is determined by the root key length you specify using the `BYTES=` parameter in the `FIELD` statement during DBD definition. If the length of the partition high key you enter is longer than the root key length, an error message displays and you must reduce the length of the partition high key. If the partition high key length is less than the defined root key length, the Partition Definition utility pads the high key value with hex 'FF's up to the defined root key length. The partition high key values must be unique for each partition within a HALDB.

The Partition High Key field consists of two sections: an editable section on the left that displays the partition high key in hexadecimal format and a view-only section on the right that displays the partition high key in alphanumeric format.

You can enter a hexadecimal value directly in the left section of the Partition High Key field. The Partition Definition utility displays the alphanumeric equivalent of this value in the right section of the Partition High Key field.

You can enter an alphanumeric value directly by using the ISPF editor. To access the ISPF editor, press F5 (If you have already entered something in the hexadecimal section, press F5 twice). Once an alphanumeric value is entered, its hexadecimal equivalent is displayed in the left section of the Partition High Key field.

An alphanumeric value can consist of any character information. If the alphanumeric value contains non-display characters, you must identify these characters using hexadecimal notation. In the ISPF editor, a hexadecimal character string is enclosed by single quotation marks and either prefixed or followed with an `x`, for example: `X'c1f201ffff'`.

Partition Selection String

The Change Partition panel displays the **Partition Selection String** field only when you have specified a partition selection routine in the HALDB master definition. A partition selection routine uses the partition selection string in hexadecimal format to distribute records across the partitions in your HALDB.

Opening HALDB Partitions

Partition selection strings are 256 bytes long. If you enter a partition selection string that is less than 256 bytes in length, the Partition Definition utility fills the remaining bytes with X'00'.

The **Partition Selection String** field consists of two sections: an editable section on the left that displays the partition selection string in hexadecimal format and a view-only section on the right that displays the partition selection string in alphanumeric format.

You can enter a hexadecimal value directly in the left section of the **Partition Selection String** field. The Partition Definition utility displays the alphanumeric equivalent of this value in the right section of the Partition Selection String field.

You can enter the partition selection string in an alphanumeric format by using the ISPF editor. To access the ISPF editor, press F5 (If you have already entered something in the hexadecimal section, press F5 twice). After you enter an alphanumeric string, its hexadecimal equivalent is displayed in the left section of the **Partition Selection String** field.

An alphanumeric string can consist of any character information. If an alphanumeric string contains non-display characters, you must identify these characters using hexadecimal notation. In the ISPF editor, a hexadecimal character string is enclosed by single quotation marks and either prefixed or followed with an x, for example: X'c1f201ffff'.

F5=String

F5 performs two functions: first, when new data is entered into the hexadecimal section of either the **Partition High Key** or the **Partition Selection String** field, F5 enters the data into the Partition Definition utility and displays the alphanumeric equivalent of the hexadecimal string in the right section of the field. Second, if there is no uncommitted data in the hexadecimal section, it displays the alphanumeric editor. Figure 317 is an example of the editor panel that is displayed for the **Partition Selection String** field.

```
EDIT          Partition Selection String

Database name . . . . : IVPDB1
Partition name . . . . : IVPD101

***** ***** Top of Data *****
=COLS> -----1-----2-----3-----4-----5---
000001 '546787789af'x
***** ***** Bottom of Data *****
Command ==>
          F1=Help  F3=Exit
```

Figure 317. Selection String Editor (DSPXPKE)

F6=Groups

Pressing F6 allows you to display the Data set group dialog that is discussed in “Defining Data Set Group Information” on page 533.

F11=List

Pressing F11 allows you to display the Database partitions panel that is discussed in “Displaying the List of Defined Partitions” on page 534.

If your definition of the HALDB from DBDLIB only allows one data set group, the **Attributes for data set group A** section is displayed. If multiple groups are

allowed, a reminder to press PF6 to work with the groups is displayed. The data set groups dialog is discussed in “Defining Data Set Group Information” on page 533.

Related Reading: For a description of the fields shown in Figure 316 on page 530, see the description for Figure 314 on page 524.

Defining Data Set Group Information

You can define data set group information by pressing F6 on the Change Partition panel (see Figure 316 on page 530). This section describes how to define the data set group information.

If you have multiple data set groups defined for your HALDB and you do not use automatic definition, use the data set group list that is displayed in Figure 318 on page 533 and Figure 319 on page 534.

From the data set groups list, you can change the attributes for each member by typing over the values in the *list* column. There is a special row in the list that allows you to make changes to an entire column of the list; the *all* row. When you type a value in the *all* row and press Enter, the value you typed is propagated to all of the members of the groups. After your changes are made, the *all* row is blanked out.

Important: Press F9 to save your changes and then press F12 to return to the previous panel.

The list contains an *action* column. The only action allowed is to display all information for a particular group. Select the group by typing a slash (/) in the *Act* column. Figure 320 on page 534 is where you can modify the values by typing over the existing data and pressing enter.

Help

Change Dataset Groups Row 1 to 11 of 11

Select an item by pressing a '/' on the desired line then press Enter.

Database name : IVPDB1
 Partition name : IVPD101
 Partition ID. : 1
 Data set name prefix. . . . : IMS.DB01.FINANCE.YEAR1998.CURR

| Act | Group | Block Size | Max Image Copies | Recovery Period | Recovery Util. JCL | Default JCL |
|-----|-------|------------|------------------|-----------------|--------------------|-------------|
| ___ | All | ___ | ___ | ___ | ___ | ___ |
| ___ | A | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | B | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | C | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | D | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | E | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | F | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | G | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | H | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | I | 8192 | 2 | 0 | RECOVJCL | ___ |
| ___ | J | 8192 | 2 | 0 | RECOVJCL | ___ |

Command ==>

F1=Help F3=Exit F7=Backward F8=Forward F9=Save F11=Right F12=Cancel

Figure 318. Change Data Set Groups, Part 1 (DSPXPGA)

Defining Data Set Group Information

Help

Change Dataset Groups

Row 1 to 11 of 11

Select an item by pressing a '/' on the desired line then press Enter.

Database name : IVPDB1
Partition name : IVPD101
Partition ID. : 1
Data set name prefix. . . . : IMS.DB01.FINANCE.YEAR1998.CURR

| Act | Group | Image Copy JCL | On. Image Copy JCL | Receive JCL | Reusable? |
|-----|-------|-------------------|-----------------------|----------------|-----------|
| ___ | All | | | | |
| ___ | A | ICJCL | OICJCL | RECVJCL | No |
| ___ | B | ICJCL | OICJCL | RECVJCL | No |
| ___ | C | ICJCL | OICJCL | RECVJCL | No |
| ___ | D | ICJCL | OICJCL | RECVJCL | No |
| ___ | E | ICJCL | OICJCL | RECVJCL | No |
| ___ | F | ICJCL | OICJCL | RECVJCL | No |
| ___ | G | ICJCL | OICJCL | RECVJCL | No |
| ___ | H | ICJCL | OICJCL | RECVJCL | No |
| ___ | I | ICJCL | OICJCL | RECVJCL | No |
| ___ | J | ICJCL | OICJCL | RECVJCL | No |

Command ==>

F1=Help F3=Exit F7=Backward F8=Forward F9=Save F11=Right F12=Cancel

Figure 319. Change Data Set Groups, Part 2 (DSPXPGB)

Help

Change a Dataset Group

Enter values, then press Enter.

Attributes for data set group B
Block size 8192

DBRC options
Max. image copies. . . . 2
Recovery period. . . . 0
Recovery utility JCL . . RECOVJCL
Default JCL.
Image copy JCL ICJCL
Online image copy JCL. . OICJCL
Receive JCL. RECVJCL
Reusable? No

Command ==>

F1=Help F3=Exit F12=Cancel

Figure 320. Change a Data Set Group (DSPXPGC)

Related Reading: For descriptions of the fields on the Change Data Set Groups panels, see the field definitions for Figure 314 on page 524.

Displaying the List of Defined Partitions

When you choose *Open database partitions* from the Partitioned Databases panel (Figure 308 on page 518), the Database Partitions list is displayed. The list is displayed immediately if there are already partitions defined for the HALDB, or it is displayed after you define partitions for HALDBs that do not already have partitions. See Figure 321 on page 535 for an example of the Database Partitions list. The list is displayed as a table that you can scroll up and down in.

File Edit View Help

Database Partitions

Row 1 to 15 of 166

Select an item by pressing a '/' on the desired line then press Enter.

Database name : IVPDB1

| Act | Name | Id | Data set name prefix | Status |
|-----|---------|----|-------------------------|----------|
| — | IVPD101 | 1 | IMS.DB01.FINANCE.YE2002 | |
| — | IVPD102 | 2 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD103 | 3 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD104 | 4 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD105 | 5 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD106 | 6 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD107 | 7 | IMS.AB01.PAYROLL.YE2002 | |
| — | IVPD108 | 8 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD109 | 9 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD110 | 10 | IMS.AB01.PAYROLL.YE2002 | Disabled |
| — | IVPD111 | 11 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD112 | 12 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD113 | 13 | IMS.TP01.PAYROLL.YE2002 | |
| — | IVPD114 | 14 | IMS.DB01.PAYROLL.YE2002 | |
| — | IVPD115 | 15 | IMS.DB01.FINANCE.YE2002 | |

Command ==>

F1=Help F3=Exit F7=Backward F8=Forward F11=Right

Figure 321. Database Partitions Panel, Sorted by Partition ID (DSPXPLA)

The Database Partitions list panel has the HALDB name at the top and table information below. Descriptions of the table columns are listed below.

Act This is the line command input field where you can invoke commands such as open, copy, and the other commands listed in “The Partition List Line Commands” on page 538.

Name The name column contains the partition name that was provided during the definition of the partition. This is the initial sort sequence.

Related Reading: For a more detailed description of the partition name, see “Opening HALDB Partitions” on page 521.

Id This is the partition ID number. The number does not have to be sequential.

Related Reading: For a more detailed description of the partition ID, see “Displaying the List of Defined Partitions” on page 534.

Data set name prefix

The data set name prefix contains the name of the data set that was provided during the definition of the partition.

Related Reading: For a more detailed description of the data set name prefix, see “Opening HALDB Partitions” on page 521.

Status A partition can be disabled by selecting a partition and typing “disable” in the Partition status field of the Change Partition panel. When a partition is disabled, “Disabled” appears in the Status column for that partition in the Database Partitions panel. For enabled partitions, the column remains blank.

From the Database Partitions list panel Figure 321, you can work with individual partitions. To use the File Action bar, type a slash (/) in the *Act* line command column for the partition you want to work with, then put the cursor on the *File* action bar choice and press Enter. Select the action you want to perform by typing the number or by positioning the cursor on the choice then pressing enter again.

Displaying Partitions List

You can invoke the Database Partitions panel (Figure 322) to show the values by pressing your PF11 key.

File Edit View Help

Database Partitions Row 1 to 4 of 166

Select an item by pressing a '/' on the desired line then press Enter.

Database name : IVPDB1

| Act | Name | Partition Selection String |
|-----|---|----------------------------|
| --- | IVPD001 | |
| | +00 F2F0F0F3 4BF2F2F4 40F1F77A F2F57AF0 | 2003.224 17:25:0 |
| | +10 F94BF6F3 F3F12432 00000000 00001020 | 9.6331..... |
| | +20 A840C1A5 85404040 40E28195 40D196A2 | y Ave San Jos |
| | +30 856B40C3 C14040F9 F5F1F4F1 00100020 | e, CA 95141.... |
| | +40 00050000 40F0F34B F0F3F440 00000100 | ... 03.034 ... |
| | +50 F1F8F0F0 C9C2D4E2 C5D9E540 40C9C2D4 | 1800IBMSERV IBM |
| | +60 40C39699 974B4040 F5F5F540 C2818993 | Corp. 555 Bail |
| | +70 A840C1A5 85404040 40E28195 40D196A2 | y Ave San Jos |
| | +80 856B40C3 C14040F9 F5F1F4F1 00403010 | e, CA 95141. .. |
| | +90 00010500 40F0F34B F2F4F340 00324020 | ... 03.243 .. |
| | +A0 9201913C D2FE933D 913C1F66 4360A005 | k.j.K.l.j....- |
| | +B0 3233A200 D996A281 6BD785A3 85996B40 | ..s.Rosa,Peter, |
| | +C0 000080D4 81A3A3F9 71C4C6F8 F1F4C6C2 | ..Matt9.DF814FB |
| | +D0 9311913C F6F4F8F6 943C1F66 4360A005 | l.j.6486m....- |
| | +E0 41E3453C 06000045 10110220 10416220 | .T..... |
| | +F0 FFFF900 00004920 18007410 94000300 | ..9.....m... |
| --- | IVPD002 | |
| | +00 F2F0F0F3 4BF2F2F4 40F1F87A F1F27AF0 | 2003.224 18:12:0 |
| | +10 F94BF6F3 F3F12432 00000000 00001020 | 9.6331..... |

Command ==>

F1=Help F3=Exit F7=Backward F8=Forward F11=Right

Figure 322. Database Partitions Panel, Sorted by Key (DSPXPLB)

The Database Partitions list panel has the HALDB name at the top and table information below. Descriptions of the table columns for Figure 322 on page 536 are presented below.

Act This is the line command input field where you can invoke commands such as open, copy, and the other commands listed on “The Partition List Line Commands” on page 538.

Name The name column contains the partition name that was provided during the definition of the partition. This is the initial sort sequence.

Related Reading: For a more detailed description of the partition name, see “Opening HALDB Partitions” on page 521.

Partition Selection String

The partition selection string is used by the partition selection routine.

Related Reading: For a more detailed description of the partition selection string, see “Opening HALDB Partitions” on page 521.

You can invoke the Database Partitions panel to show the Randomizer values by pressing your PF11 key (Figure 323 on page 537).

```

File Edit View Help
-----
Database Partitions                               Row 1 to 15 of 166

Select an item by pressing a '/' on the desired line then press Enter.

Database name . . . . . : IVPDB1

----- Randomizer -----
- Free Space -
Act  Name      Module    Anchor  High block  Bytes    FBFF    FSPF
----  -
IVPD101 DD41DUP2    2       999       2000      0       0
IVPD102 DD41DUP2    2       999       2000      0       0
IVPD103 DD41DUP2    2       999       2000      0       0
IVPD104 DD41DUP2    2       999       2000      0       0
IVPD105 DD41DUP2    2       999       2000      0       0
IVPD106 DD41DUP2    2       999       2000      0       0
IVPD107 DD41DUP2    2       999       2000      0       0
IVPD108 DD41DUP2    2       999       2000      0       0
IVPD109 DD41DUP2    2       999       2000      0       0
IVPD110 DD41DUP2    2       999       2000      0       0
IVPD111 DD41DUP2    2       999       2000      0       0
IVPD112 DD41DUP2    2       999       2000      0       0
IVPD113 DD41DUP2    2       999       2000      0       0
IVPD114 DD41DUP2    2       999       2000      0       0
IVPD115 DD41DUP2    2       999       2000      0       0
-----

Command ==>
F1=Help  F3=Exit  F7=Backward  F8=Forward  F11=Right

```

Figure 323. Database Partitions Panel, Sorted by Name (DSPXPLC)

Important: The Randomizer section is present only if the HALDB is PHDAM.

The Database Partitions list panel has the HALDB name at the top and table information below. Descriptions of the table columns for Figure 323 on page 537 are presented below.

Act This is the line command input field where you can invoke commands such as open, copy, and the other commands listed in “The Partition List Line Commands” on page 538.

Name The name column contains the partition name provided during the definition of the partition. This is the initial sort sequence.

Related Reading: For a more detailed description of the partition name, see “Opening HALDB Partitions” on page 521.

Module

The module column contains the module name of the randomizing module.

Related Reading: For a more detailed description of the module name see Figure 314 on page 524.

Anchor

The anchor column contains the number of root anchor points.

Related Reading: For a more detailed description of the anchor see Figure 314 on page 524.

High block

The high block column contains the high block number.

Related Reading: For a more detailed description of the high block number see Figure 314 on page 524.

Bytes For a more detailed description of the bytes field see Figure 314 on page 524.

FBFF The FBFF column contains the free block frequency factor.

Displaying Partitions List

Related Reading: For a more detailed description of the free block frequency factor in Figure 314 on page 524.

FSPF The FSPF column contains the free space percentage factor.

Related Reading: For a more detailed description of the free space percentage factor see Figure 314 on page 524.

To use line commands, type the command in the *Act* column to the right of the partition you want to use. You can type multiple line commands (only one per partition, though) on the Database Partitions panel: the commands are executed serially starting from the top.

The Partition List Line Commands

Line commands will allow you to perform the following actions:

Delete a partition

Type a D in the line command field and press Enter. A delete confirmation panel is displayed. Type a 1 in the option field and press Enter to confirm the delete or press the cancel key to cancel the delete.

If you are deleting several partitions at once, and wish to accept all of the deletes, you can type a 2 in the option field. It is reset to blank each time the partition list is displayed.

Copy a partition

Type a C in the line command field to define a new partition using the attributes of the selected partition. The partition name must be unique. You change the partition information using the Change Partition panel (see Figure 316 on page 530).

Open a partition

Type an O in the line command field to open a partition. You can then change partition information. The partition name and ID cannot be changed. press Enter to commit or press the cancel key to discard your changes. You change the partition information using the Change Partition panel (see Figure 316 on page 530).

Print partition information

Type a P in the line command field to print partition information for the selected partition. The information will not be routed to a printer immediately; instead it is added to the ISPF list data set.

The Partition List Action Bar

The list of partitions in the Database Partitions panel can be manipulated with line commands or by using the File action bar choices (Figure 324).

File Edit View Help

1. New partition . . .
2. Open partition
3. Open dataset groups
4. Print partition information
5. Print partition view

Figure 324. File Action Bar Choices

New partition

You can create new partitions using the same panels that you used when you initially created partitions. See Figure 314 on page 524.

Open partition

You can open the selected partitions and modify them as desired. See Figure 316 on page 530.

Open data set groups

You can manipulate the data set group members using the panels described in “Defining Data Set Group Information” on page 533.

Print partition information

Information about the selected partitions is written to the ISPF list data set.

Print partition view

The information in the currently-displayed view is written to the ISPF list data set.

The list of partitions in the Database Partitions can be sorted in various ways using the Edit action bar choice (Figure 325).

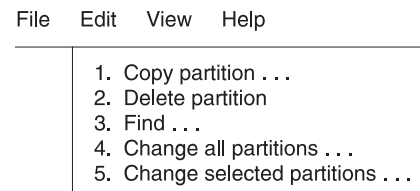


Figure 325. Edit Action Bar Choices

Copy partition...

Type a slash (/) in the line command field and use the Edit - Copy partition pull-down panel to define a new partition using the attributes of the selected partition. The partition name and the ID must be unique.

The Change Partition panel is then displayed, see Figure 316 on page 530, and you can create new partitions serially. The values shown in the panel are filled-in using the attributes of the selected partition.

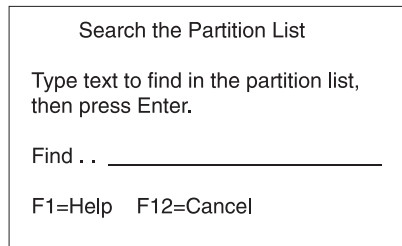
Delete partition

Type a slash (/) in the line command field and use the Edit - Delete a partition pull-down panel to delete partitions. A delete confirmation panel is displayed. You can press Enter to confirm delete or press the cancel key to ignore the delete.

Find... You can search the partition list for a selected character string. Only simple character values can be specified. The cursor is placed on the partition that contains the search value.

The search string is not case sensitive. It will search on any field, not just the currently displayed fields on the Database Partitions panels (Figure 326 on page 540).

Displaying Partitions List



Search the Partition List

Type text to find in the partition list,
then press Enter.

Find . . _____

F1=Help F12=Cancel

Figure 326. Searching the Partition List

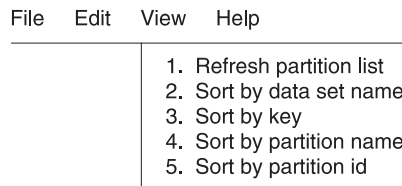
Change all partitions...

Use the Edit - Change all partitions pull-down panel to change individual fields for all of the partitions in the HALDB. The partition name and the ID cannot be changed. See “Change All Partitions.”

Change selected partitions...

Type a slash (/) in the line command field to change a partition and use the Edit - Change selected partitions pull-down panel to change individual fields for the selected partitions. The partition name and the ID cannot be changed. The process is the same as that described in “Change All Partitions,” but only the selected partitions are changed.

The list of partitions in the Database Partitions can be sorted in various ways by using the View action bar choice (Figure 327).



File Edit View Help

- 1. Refresh partition list
- 2. Sort by data set name
- 3. Sort by key
- 4. Sort by partition name
- 5. Sort by partition id

Figure 327. View Action Bar Choices

Help information is available using the Help action bar choice.

Change All Partitions

Figure 328 on page 541 is an example of the Change Partition panel. The entry fields are blank. Make changes only to the fields that you want to change. The field changes are applied to all of the partitions.

Important: The same process is used for Change selected partitions except that the changes are only applied to the partitions selected from the list with a slash (/).

If you want to change a character field to blanks, type a single slash (/) character so that it is the only character in the field.

Help

Change Partition

Press Enter to continue.

Database name. : IVPDB1
Partition name :
Partition ID :
Data set name prefix _____
Status. _____

Partition Selection String

| | | | | |
|-----|-------|-------|-------|-------|
| +00 | _____ | _____ | _____ | _____ |
| +10 | _____ | _____ | _____ | _____ |
| +20 | _____ | _____ | _____ | _____ |
| +30 | _____ | _____ | _____ | _____ |
| +40 | _____ | _____ | _____ | _____ |
| +50 | _____ | _____ | _____ | _____ |
| +60 | _____ | _____ | _____ | _____ |
| +70 | _____ | _____ | _____ | _____ |
| +80 | _____ | _____ | _____ | _____ |
| +90 | _____ | _____ | _____ | _____ |
| +A0 | _____ | _____ | _____ | _____ |
| +B0 | _____ | _____ | _____ | _____ |
| +C0 | _____ | _____ | _____ | _____ |
| +D0 | _____ | _____ | _____ | _____ |
| +E0 | _____ | _____ | _____ | _____ |
| +F0 | _____ | _____ | _____ | _____ |

Randomizer
Module name DD41MOD3
Anchor _____
High block number. _____
Bytes _____

Free Space
Free block freq. factor. . _____
Free space percentage. . . _____

Command ==>
F1=Help F3=Exit F5=String F6=Groups F12=Cancel

Figure 328. Change Partition Panel (DSPXPPB)

Important:

- The Randomizer section is present only if the HALDB is PHDAM.
- The data set groups section is present only if there is only one data set group specified during DBDGEN. If there is more than one data set group, use F6=Groups to display all data set groups using the dialog described in “Defining Data Set Group Information” on page 533.

Figure 329 on page 542 shows the Change Dataset Groups panel.

Displaying Partitions List

```
Help
-----
Change Dataset Groups          Row 1 to 10 of 10

Select an item by pressing a '/' on the desired line then press Enter.

Database name . . . . . : IVPDB1
Partition name . . . . . : IVPD101
Partition ID. . . . . : 1
Data set name prefix. . . : IMS.DB01.FINANCE.YEAR1998.CURR

Act  Group  Block  Max Image  Recovery  Recovery  Default
   Size    Copies  Period    Util. JCL  JCL
---  ---    ---    ---        ---        ---    ---
___  All    ___    ___        ___        ___    ___
___  A      ___    ___        ___        ___    ___
___  B      ___    ___        ___        ___    ___
___  C      ___    ___        ___        ___    ___
___  D      ___    ___        ___        ___    ___
___  E      ___    ___        ___        ___    ___
___  F      ___    ___        ___        ___    ___
___  G      ___    ___        ___        ___    ___
___  H      ___    ___        ___        ___    ___
___  I      ___    ___        ___        ___    ___
___  J      ___    ___        ___        ___    ___
Command ==>
F1=Help  F3=Exit  F7=Backward  F8=Forward  F9=Save  F11=Right  F12=Cancel
```

Figure 329. Change Data Set Groups, Part 1 (DSPXPGA)

Related Reading: For a description of the fields not listed here, see the description for Figure 313 on page 522.

Opening Database Information

When you choose Open database information from the Partitioned Databases panel shown in Figure 308 on page 518, you are shown information about the HALDB which was saved when you first defined partitions for the HALDB (Figure 330).

```
Help
-----
Partitioned Database Information

Type the field values. Then press Enter to continue.

Database name . . . . . : IVPDB1

Master Database values
Part. selection routine . . . DFSIVD1
RSR global service group . . .
RSR tracking type . . . . .
Share level . . . . . 0
Database organization . . . : PHDAM
Recoverable? . . . . . Yes
Number of data set groups . : 1
Online Reorganization Capable: Yes

To exit the application, press F3.

Command ==>
F1=Help  F3=Exit  F12=Cancel
```

Figure 330. Partitioned Database Information (DSPXPOA)

Related Reading: For a description of the fields shown in Figure 330, see the description for Figure 313 on page 522.

You can modify the fields and press Enter to change the values in RECON. If you press cancel or exit, any changes you entered on this panel are discarded.

Deleting Database Information

When you choose to delete database information from Figure 308 on page 518, you are presented with the Delete a Database panel (see Figure 331). You must type a slash (/) character and press Enter to confirm the delete. When you confirm it, the information about the HALDB and about all of its partitions is deleted from RECON.

There is no way to undo the delete. You may wish to perform an export prior to deleting a HALDB from RECON. See “Exporting Database Information” for information about performing an export.

```

Help
-----
                        Delete Database Information

Type '/' to confirm the delete of the database information from RECON.
Then press Enter.

Database name . . . . . : IVPDB1

Confirm database delete . __

Command ==>
    F1=Help    F3=Exit    F12=Cancel

```

Figure 331. Delete a Database (DSPXPDA)

Exporting Database Information

When you choose to Export database information from Figure 308 on page 518, the information is stored in the partitioned data set that you specify. It is saved as an ISPF table and so must have the attributes of ISPTLIB data sets (record format = fixed block, record length = 80, data set organization = PDS (or PDS/E)). Figure 332 shows the Export a Database panel.

```

Help
-----
                        Export a Database

Type a data set name. Then press Enter.

Database Name . . . . . : IVPDB1
Output dataset name. . . . 'TEST.RSR.PARTS'
Output member name . . . . IVPDB1

To exit the application, press F3.

Command ==>
    F1=Help    F12=Cancel

```

Figure 332. Export a Database (DSPXPEA)

Field Description

Database name

The HALDB name that was specified in the primary panel.

Exporting Database Information

Output data set name

The output data set name is the name of the data set that will contain the partition information.

Importing Database Information

When you choose to Import database information from Figure 308 on page 518 you can specify the name of the PDS or PDS/E that contains the information.

Important: Only an exported table can be used for the import.

After you press Enter, the table is read and each partition is defined.

Figure 333 shows the Import a Database panel.

```
Help
-----
                        Import a Database

Type a dataset name. Then press Enter.

Database name . . . . : IVPDB1
Input dataset name. . . . 'PROD.RSR.PARTS'
Input member name . . . . IVPDB1

Processing option . . . . _ 1. Stop on first error
                        2. Try all partitions

Command ==>
F1=Help  F3=Exit  F12=Cancel
```

Figure 333. Import a Database (DSPXP1A)

Database name

The HALDB name that was specified on the primary panel.

Input data set name

The input data set name is the name of the data set that contains the partition information. The data set must be partitioned.

Input member name

The input member name is the name of a member within the input data set. The member must have been exported using the HALDB Partition Definition utility .

Processing option

Each partition in the imported table can be defined in RECON. If there are errors, you can choose to try the remaining partitions or to stop the process.

Displaying the IMS Concatenation

You can look at the concatenation of data sets that are allocated to the IMS DD name. The data sets are displayed using the ISRDDN command that is part of the ISPF product (Figure 334 on page 545).

| Current Data Set Allocations | | | | | | Line 1 of 2 | |
|------------------------------------|-------------|-----|--------|----------------------|------|------------------------|--|
| Volume | Disposition | Act | DDname | Data Set Name | List | Actions: B E V F C I Q | |
| SYS151 | SHR,KEEP | > _ | IMS | IMSIVP91.DBDLIB | | | |
| SYS335 | SHR,KEEP | > _ | | IMS91.SANJOSE.DBDLIB | | | |
| ----- End of Allocation List ----- | | | | | | | |
| Command ==> | | | | | | Scroll ==> CSR | |

Figure 334. The IMS Concatenation (ISRDDN)

Use the help (F1) information provided by ISRDDN and in the ISPF manuals to learn more about the ISRDDN utility. When you exit the ISRDDN utility, the HALDB Partition Definition utility panels are displayed again.

Selecting an IMS Configuration Using the Partition Definition Utility

You can control which RECON data set and which DBDLIB data sets are used. For this purpose, a set of RECON and DBDLIB data sets are considered a configuration.

The configuration is a name that you specify that identifies a set of DBD libraries and a set of RECON data sets. If you already have the IMS DD name allocated from the logon procedure and if you have the IMS.SDFSRESLs allocated to the STEPLIB DD name, you do not need to use the Configuration option. If you do define and select a configuration, those data sets will override the allocations from the logon procedure.

1. IMS DD name

The IMS DD name includes the data sets that contain the DBDLIB members. The RECON / DBDLIB Configurations panels re-allocate the IMS DD name.

2. RECON allocation

The STEPLIB allocation contains RECON1, RECON2, and RECON3 members that name the actual RECON data sets. IMS uses those members to determine which RECON data sets to use. There is an alternative to using a STEPLIB: use the TSOLIB command to change the search order that TSO/E uses to find commands and programs.

The RECON / DBDLIB Configurations panels re-allocates the IMS DD name and will allocate the RECON1, RECON2, and RECON3 DDnames to explicitly specify the RECON data sets. The STEPLIB concatenation is not modified.

Figure 335 on page 546 shows the RECON/DBDLIB Configurations panel.

Selecting an IMS Configuration Using the Partition Definition Utility

```

                                RECON / DBDLIB Configurations                                Row 1 to 5 of 5

To create a new configuration, fill in the first line and press Enter.
Select a default by type '/' on the Act column then press Enter.
You can use 'O' to open or 'D' to delete a configuration.

Act   Current  Name           Description
-----
----- *      SDFSRESL      IMS V9R1 datasets
-----          TESTM        Test IMS for Matt
-----          TESTP        Test IMS for Peter
-----          TESTI        Test IMS
***** Bottom of data *****

Command ==>
F1=Help      F3=Exit      F7=Up        F8=Down      F12=Cancel

```

Figure 335. User Configurations (DSPXPMB)

A list of configurations can be maintained when you select option 7 from the Partitioned Databases panel. The list is initially empty and it can be added-to by filling in the blank line. The active configuration is identified by an asterisk (*) in the Current column. Figure 336 shows the Configurations Details panel.

Rows from the list can be deleted by using a line command of d. Only the configuration is deleted from the list. The data sets that are named in the configuration are not deleted.

The data sets named in the configuration are set or changed by using a line command of o for open.

```

                                Configuration Details

Type in values in the fields and press Enter to continue.

Configuration name . . . . TEST1

Description . . . . . Test IMS

RECON dataset names
  RECON1 dataset . . . . 'TEST.PARTS.RECON1'
  RECON2 dataset . . . . 'TEST.PARTS.RECON2'
  RECON3 dataset . . . . 
DBDLIB dataset names
  DBDLIB dataset 1 . . . 'TEST.PARTS.DBDLIB'
  DBDLIB dataset 2 . . . 
  DBDLIB dataset 3 . . . 
  DBDLIB dataset 4 . . . 
  DBDLIB dataset 5 . . . 
  DBDLIB dataset 6 . . . 
  DBDLIB dataset 7 . . . 
  DBDLIB dataset 8 . . . 
  DBDLIB dataset 9 . . . 
  DBDLIB dataset 10 . . . 

Command ==>
F13=Help     F15=Exit     F19=Up       F20=Down     F22=Actions

```

Figure 336. Configuration Details Panel (DSPXPMC)

Selecting an IMS Configuration Using the Partition Definition Utility

The RECON data sets are separately allocated to the RECON1, RECON2, and RECON3 file names.

The DBDLIB data sets are concatenated to the IMS file name.

Important: When you specify a generic HALDB name in the Partitioned Database panel; option 6 will only work if you use four (4) or fewer DBD data sets. However, for greater flexibility you can specify up to ten (10) data sets.

Using Batch to Export or Import Partition Information

The output from the export of a HALDB is a member of a PDS. The information about the HALDB is saved in the form of an ISPF table. The ISPF table is used as input for the import process. The import can be done from the ISPF panels or from a batch job.

The batch import of a HALDB can be done by submitting a batch ISPF job similar to the job shown in Figure 337. ISPF is invoked in batch, so all ISPF DDNAMES are required.

```
//DSPXRUN JOB ...
//*
//DSPXRUN EXEC PGM=IKJEFT01,DYNAMNBR=50,REGION=6M
//STEPLIB DD DSN=IMSIVP91.SDFSRESL,DISP=SHR /* IMS.SDFSRESL */
//SYSPROC DD DSN=IMSIVP91.SDFSEEXEC,DISP=SHR /* IMS rexx execs */
//IMS DD DSN=your.local.DBDLIB,DISP=SHR
//RECON1 DD DSN=IMSIVP91.RECON1,DISP=SHR
//RECON2 DD DSN=IMSIVP91.RECON2,DISP=SHR
//RECON3 DD DSN=IMSIVP91.RECON3,DISP=SHR
//ISPPROF DD DSN=&&PROFILE;, /* dummy ISPF profile */
// UNIT=SYSDA,DISP=(NEW,DELETE),
// SPACE=(3200,(30,30,1)),DCB=(RECFM=FB,LRECL=80,BLKSIZE=3200)
//ISPLIB DD DSN=IMSIVP91.SDFSPLIB,DISP=SHR /* IMS ISPF panels */
//ISPSLIB DD DSN=IMSIVP91.SDFSMLIB,DISP=SHR /* IMS ISPF skeletons */
//ISPLMLIB DD DSN=IMSIVP91.SDFSMLIB,DISP=SHR /* IMS ISPF messages */
// DD DSN=ISP.ISPLMLIB,DISP=SHR
//ISPTLIB DD DSN=IMSIVP91.SDFSSTLIB,DISP=SHR /* IMS ISPF tables */
// DD DSN=ISP.ISPTLIB,DISP=SHR
//ISPLIB DD SYSOUT=*,DCB=(RECFM=VA,LRECL=125,BLKSIZE=129)
//SYSPRINT DD SYSOUT=*,DCB=(RECFM=VA,LRECL=125,BLKSIZE=129)
//SYSOUT DD SYSOUT=*
//PARTLOG DD SYSOUT=*
//SYSTSPRT DD SYSOUT=*,DCB=(RECFM=F,LRECL=255,BLKSIZE=255)
//SYSTSIN DD *
ISPSTART CMD( +
DSPXRUN IMPORT DSN('PROD.RSR.PARTS') +
DBN(IVPDB1) MEM(IVPDB1) OPT(2))
/*
```

Figure 337. Sample JCL for Batch Import

The batch job executes the standard ISPF command ISPSTART that sets up the ISPF environment then starts the DSPXRUN command. The DSPXRUN command identifies the HALDB, the import file to use, and the processing options.

DSPXRUN Command Syntax

The DSPXRUN command can be used to import HALDB information in a batch environment.

The following diagram shows the syntax of the DSPXRUN command:

►—ISPSTART—CMD—(| Command String |)—

Command String:

|—DSPXRUN—
| EXPORT |
| IMPORT |—DBN—(database_name)—DSN—(dataset_name)—

►—MEM—(member_name)—OPT—(processing_option)—

The values are essentially the same as the values required for the foreground import (see “Importing Database Information” on page 544).

EXPORT

When you choose to export database information using a batch job, the information is stored in the partitioned data set that you specify. The information is saved as an ISPF table and so it must have the attributes of ISPTLIB data sets: record format = fixed block, record length = 80, and data set organization = PDS (or PDS/E).

Related Reading: For more information on ISPTLIB data sets, see *ISPF User's Guide, Volume 1*.

IMPORT

When you choose to import database information using a batch job, the partition information is read from a partitioned data set that you specify. The partition information is defined to the RECON data sets.

database_name

The HALDB name that was specified on the primary panel.

dataset_name

The input data set name is the name of the data set that contains the partition information. The data set must be partitioned.

member_name

The input member name is the name of a member within the input data set. The member must have been exported using the HALDB Partition Definition utility.

processing_option

The processing option field lets you determine what the Partition Definition utility does in the event that an error occurs when it processes a partition from the imported table. The Partition Definition utility records each partition it imports in RECON. If there are errors, you can choose to try the remaining partitions or to stop the process. The valid values are 1 or 2:

- 1 Stop on first error (prior imported partitions are retained)
- 2 Try all partitions

The OPT parameter is ignored during export processing.

DSPXRUN EXPORT Sample Output

Exporting database IVPDB1 using JCL similar to Figure 337 on page 547 would result in the following output:

```
DSPM142I Start export to MEM=IVPDB1 in DSN='PROD.RSR.PARTS'
        from DBN=IVPDB1
DSPM143I The export file contains partition IVPDB11
DSPM143I The export file contains partition IVPDB12
DSPM143I The export file contains partition IVPDB13
DSPM219I Table IVPDB1 was created successfully to dataset
        'PROD.RSR.PARTS'
```

DSPXRUN IMPORT Sample Output

Importing database IVPDB1 using JCL similar to Figure 337 on page 547 would result in the following output:

```
DSPM283I Start Import to DBN=IVPDB1 from MEM=IVPDB1 in
        DSN='PROD.RSR.PARTS' Options=2
DSPM285I Imports start at 23/22/14 11:55
DSPM284I Import successful for partition IVPDB11
DSPM284I Import successful for partition IVPDB12
DSPM284I Import successful for partition IVPDB13
DSPM282I 3 of a total 3 partitions from table IVPDB1
        were imported to database successfully.
```

Appendix F. Output Data Set Requirements for HALDB Online Reorganization

This appendix describes the detailed requirements and attributes for the output data sets used during online reorganization of HALDB partitions. These requirements are described in the following topics:

- “HALDB Online Reorganization Requirements for Existing Output Data Sets”
- “Attributes of Automatically-Created Output Data Sets”

HALDB Online Reorganization Requirements for Existing Output Data Sets

If an existing output data set does not meet the requirements described in this section, IMS displays an error message and the online reorganization for the HALDB partition does not begin.

An OSAM output data set has the following requirements:

- Must be cataloged
- Must be a DASD data set
- Must not be a VSAM data set, except for the primary index data set of a PHIDAM database
- Must not be a PDS, PDSE, or a member of a PDS or PDSE

A VSAM output data set has the following requirements:

- Must be a VSAM entry-sequenced data set (ESDS), except for the primary index data set of a PHIDAM database
- Must have the REUSE attribute
- Must have a fixed-length record length that is identical to that of the corresponding input data set
- Must have a control interval size that is identical to that of the corresponding input data set
- Must have a SHAREOPTIONS attribute value that is at least as high as that of the corresponding input data set if the database is defined to DBRC with a SHARELVL attribute value of 2 or 3

A primary index data set has the following requirements:

- Must be a VSAM key-sequenced data set (KSDS)
- Must have the same key offset and length as the corresponding input KSDS
- Must have the other required characteristics listed for VSAM output data sets

Attributes of Automatically-Created Output Data Sets

For those output data sets that do not already exist at the beginning of the initialization phase of an online reorganization, IMS creates the data sets with the following attributes:

Number of Volumes

If a particular input data set is SMS-managed, IMS creates the corresponding output data set with the same number of volumes.

If the input data set is not SMS-managed, IMS automatically creates the corresponding output data set only when the input data set resides on a single volume. For a non-SMS-managed input data set that resides on multiple volumes, you must create the corresponding output data set before starting the online reorganization.

Location of SMS-managed output data sets

If a particular input data set is SMS-managed, the corresponding output data set is also SMS-managed, and uses the same storage class as the input data set.

Your site's storage administrator must ensure that this storage class refers to a storage group with sufficient space to hold the output data set, or that the automatic class selection (ACS) routine selects an appropriate storage class for the data set.

Location of non-SMS-managed, non-VSAM output data sets

Regardless of the type of DASD on which the input data set resides, IMS creates the corresponding non-VSAM output data set using the equivalent of a DD statement UNIT=SYSALLDA parameter.

When it creates the output data set, IMS does not request any specific volume serial number, thus allowing the data set to be created on a storage volume or, if no storage volume is available, on a public volume.

Location of non-SMS-managed, VSAM output data sets

IMS creates a VSAM output data set on the same volume as the corresponding input data set. This restriction can limit the usefulness of automatically creating a VSAM data set that is not SMS-managed.

Size of output data sets on a single volume

When the input data set has extents on only one DASD volume, IMS creates the output data set on a single volume using the equivalent of a DD statement VOLUME=(,,1) parameter.

The amount of primary space for the output data set is derived from the space allocation of the input data set:

- For a non-VSAM data set, the primary space is the total amount of space in the first five extents on the volume.
- For a VSAM data set, the primary space is the primary space allocation used when the input data set was created.

If you specified secondary space amount for the input data set, IMS uses this same secondary amount for the output data set.

To reserve approximately the same amount of space for the output data set as was reserved for the input data set, regardless of the DASD types involved, IMS requests the space for the output data set as a number of OSAM blocks or VSAM records. For input data sets that did not specify a number of OSAM blocks or VSAM records, IMS converts the cylinder or track allocation to an equivalent number of blocks or records.

An automatically created output data set could have a considerably different amount of available DASD space than was used for the input data set. For example, for an input data set that used

secondary allocation, the automatic creation process reserves the primary space for the output data set, but there might not be enough space on the volume for secondary allocation either during the online reorganization or during later database processing.

Size of output data sets on multiple volumes (SMS-managed only)

IMS automatically creates multiple-volume output data sets only when the input data set (and, therefore, the output data set) is SMS-managed. You can determine the storage class by examining the input data set or the site's ACS routine.

Although it is not strictly a requirement for SMS-managed multiple-volume output data sets, you should ensure that the storage class specifies the guaranteed-space attribute. By specifying the guaranteed-space attribute, you allow VSAM to use the primary-space allocation for *each* of the volumes when it creates the output data sets. Secondary space is used as needed. However, even with the guaranteed-space attribute, the output data sets might not have the same amount of space as the input data sets, especially if secondary-space allocation was used for the input data sets.

The requested primary and secondary space is based on the input data set's space allocation on the first DASD volume.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs

and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming Interface Information

This book is intended to help the database administrator manage IMS databases.

This book also documents general-use interface and Associated Guidance Information, Product-sensitive Programming Interface and Associated Guidance Information, and diagnosis, modification or tuning Information provided by IMS.

General-use programming interfaces allow the customer to write programs that obtain the services of IMS.

General-use Programming Interface and Associated Guidance Information is identified where it occurs by an introductory statement to a chapter or section.

Product-sensitive programming interfaces allow the customer installation to perform tasks such as diagnosing, modifying, monitoring, repairing, tailoring, or tuning of IMS. Use of such interfaces creates dependencies on the detailed design or implementation of the IBM software product. Product-sensitive programming interfaces should be used only for these specialized purposes. Because of their dependencies on detailed design and implementation, it is to be expected that programs written to such interfaces may need to be changed in order to run with new product releases or versions, or as a result of service.

Product-sensitive Programming Interface and Associated Guidance Information is identified where it occurs by an introductory statement to a chapter or section.

Diagnosis, modification or tuning information is provided to help the customer diagnose, modify, or tune IMS.

Attention: Do not use this Diagnosis, Modification or Tuning Information as a programming interface.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States, other countries, or both:

| | |
|------------------------|-----------|
| BookManager | NetView |
| CICS | OS/390 |
| DataPropagator | QMF |
| DataRefresher | RACF |
| DB2 | RAMAC |
| DB2 Universal Database | Redbooks |
| DFSMSdss | RMF |
| Hiperspace | SAA |
| IBM | Tivoli |
| IMS | WebSphere |
| MVS | z/OS |

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

UNIX is a trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Bibliography

This bibliography lists all of the information in the IMS Version 9 library.

CICS Transaction Server for z/OS V2.3: CICS RACF Security Guide, SC34-6249

Data Extraction, Processing, and Restructuring System Program Description/Operations Manual, SH20-2177

DB2 Universal Database for z/OS V8: DB2 Universal Database for z/OS Administration Guide, SC26-9931

File Manager for z/OS User's Guide and Reference for IMS Data, SC18-9382.

IBM DB2 and IMS Tools: IMS Application Development Facility (ADF) II User's Guide, SH20-6593

IBM DB2 and IMS Tools: IMS DataPropagator for z/OS: An Introduction, GC27-1211

IBM DB2 and IMS Tools: IMS High Availability Large Database Conversion and Maintenance Aid for z/OS, User's Guide, SC18-7249

IBM DB2 and IMS Tools: IMS High Performance Image Copy for z/OS User's Guide, SC18-7617

IBM DB2 and IMS Tools: IMS High Performance Load for OS/390 User's Guide, SC27-0938

IBM DB2 and IMS Tools: IBM IMS High Performance Pointer Checker for z/OS User's Guide Volume 2, SC18-7258

IBM DB2 and IMS Tools: IMS High Performance Unload for OS/390, SC27-0936

IBM DB2 and IMS Tools: IMS Index Builder for z/OS User's Guide, SC27-0930

IBM Redbooks: The Complete IMS HALDB Guide, SG24-6945

VisualAge Generator: VisualAge Generator User's Guide, SH23-0268-01

z/OS V1R4: DFSMS Access Method Services for Catalogs, SC26-7394

z/OS V1R4: DFSMS: Using Data Sets, SC26-7410

z/OS V1R4: DFSMSdss Storage Administration Reference, SC35-0424

z/OS V1R4: MVS Setting Up a Sysplex, SA22-7625

z/OS V1R4: MVS System Messages, Vol 7 (IEB-IEE), SA22-7637

IMS Version 9 Library

| Title | Acronym | Order number |
|---|---------|--------------|
| <i>IMS Version 9: Administration Guide: Database Manager</i> | ADB | SC18-7806 |
| <i>IMS Version 9: Administration Guide: System</i> | AS | SC18-7807 |
| <i>IMS Version 9: Administration Guide: Transaction Manager</i> | ATM | SC18-7808 |
| <i>IMS Version 9: Application Programming: Database Manager</i> | APDB | SC18-7809 |
| <i>IMS Version 9: Application Programming: Design Guide</i> | APDG | SC18-7810 |
| <i>IMS Version 9: Application Programming: EXEC DLI Commands for CICS and IMS</i> | APCICS | SC18-7811 |
| <i>IMS Version 9: Application Programming: Transaction Manager</i> | APTM | SC18-7812 |
| <i>IMS Version 9: Base Primitive Environment Guide and Reference</i> | BPE | SC18-7813 |
| <i>IMS Version 9: Command Reference</i> | CR | SC18-7814 |
| <i>IMS Version 9: Common Queue Server Guide and Reference</i> | CQS | SC18-7815 |
| <i>IMS Version 9: Common Service Layer Guide and Reference</i> | CSL | SC18-7816 |
| <i>IMS Version 9: Customization Guide</i> | CG | SC18-7817 |
| <i>IMS Version 9: Database Recovery Control (DBRC) Guide and Reference</i> | DBRC | SC18-7818 |
| <i>IMS Version 9: Diagnosis Guide and Reference</i> | DGR | LY37-3203 |
| <i>IMS Version 9: Failure Analysis Structure Tables (FAST) for Dump Analysis</i> | FAST | LY37-3204 |
| <i>IMS Version 9: IMS Connect Guide and Reference</i> | CT | SC18-9287 |
| <i>IMS Version 9: IMS Java Guide and Reference</i> | JGR | SC18-7821 |
| <i>IMS Version 9: Installation Volume 1: Installation Verification</i> | IIV | GC18-7822 |
| <i>IMS Version 9: Installation Volume 2: System Definition and Tailoring</i> | ISDT | GC18-7823 |
| <i>IMS Version 9: Master Index and Glossary</i> | MIG | SC18-7826 |

| Title | Acronym | Order number |
|---|----------------|---------------------|
| <i>IMS Version 9: Messages and Codes, Volume 1</i> | MC1 | GC18-7827 |
| <i>IMS Version 9: Messages and Codes, Volume 2</i> | MC2 | GC18-7828 |
| <i>IMS Version 9: Open Transaction Manager Access Guide and Reference</i> | OTMA | SC18-7829 |
| <i>IMS Version 9: Operations Guide</i> | OG | SC18-7830 |
| <i>IMS Version 9: Release Planning Guide</i> | RPG | GC17-7831 |
| <i>IMS Version 9: Summary of Operator Commands</i> | SOC | SC18-7832 |
| <i>IMS Version 9: Utilities Reference: Database and Transaction Manager</i> | URDBTM | SC18-7833 |
| <i>IMS Version 9: Utilities Reference: System</i> | URS | SC18-7834 |

Accessibility Titles Cited in This Library

| Title | Order number |
|--|---------------------|
| <i>z/OS V1R1.0 TSO Primer</i> | SA22-7787 |
| <i>z/OS V1R5.0 TSO/E User's Guide</i> | SA22-7794 |
| <i>z/OS V1R5.0 ISPF User's Guide, Volume 1</i> | SC34-4822 |

Supplementary Publications

| Title | Order number |
|---|---------------------|
| <i>IMS Connector for Java 2.2.2 and 9.1.0.1 Online Documentation for WebSphere Studio Application Developer Integration Edition 5.1.1</i> | SC09-7869 |
| <i>IMS Version 9 Fact Sheet</i> | GC18-7697 |
| <i>IMS Version 9: Licensed Program Specifications</i> | GC18-7825 |

Publication Collections

| Title | Format | Order number |
|---|-----------------|---------------------|
| <i>IMS Version 9 Softcopy Library</i> | CD | LK3T-7213 |
| <i>IMS Favorites</i> | CD | LK3T-7144 |
| <i>Licensed Bill of Forms (LBOF): IMS Version 9 Hardcopy and Softcopy Library</i> | Hardcopy and CD | LBOF-7789 |
| <i>Unlicensed Bill of Forms (SBOF): IMS Version 9 Unlicensed Hardcopy Library</i> | Hardcopy | SBOF-7790 |
| <i>OS/390 Collection</i> | CD | SK2T-6700 |
| <i>z/OS Software Products Collection</i> | CD | SK3T-4270 |
| <i>z/OS and Software Products DVD Collection</i> | DVD | SK3T-4271 |

Index

Special characters

/CK operand 202
/START AREA usage 459, 461
/START DATABASE command 459
/SX operand 202

A

abnormal termination in logical relationships 503, 505
ACB (application control block)
 description 308
 generation 308
ACBGEN (Application Control Block Generation)
 utility 459
ACBLIB
 See IMS.ACBLIB library
access methods
 BSAM (Basic Sequential Access Method) 513
 converting DL/I access methods 394
 DL/I access methods, changing 394
 HISAM 69
 IMS access methods 11, 64
 introduction 11
 operating system access methods 11
 OSAM (Overflow Sequential Access Method) 513
 OSAM (overflow sequential access methods)
 used by HD 95
 QSAM (Queued Sequential Access Method) 513
 VSAM
 HISAM 69
 z/OS access methods
 used by HD 83
 used by HSAM 65
accessing segments
 HDAM (Hierarchical Direct Access Method) 103
 HIDAM (Hierarchical Indexed Direct Access Method) 103
 HISAM (Hierarchical Indexed Sequential Access Method) 72
 HSAM (Hierarchical Sequential Access Method) 67
 PHDAM (Partitioned Hierarchical Direct Access Method) 103
 PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) 103
add programs, use in loading a database 326
adjusting HDAM options 247
adjusting PHDAM options 247
administration
 database
 overview 3
aids
 for test databases
 Data Extraction, Processing and Restructuring System 313
 DL/I test program 314
 VisualAge Generator 313
AL (available length) field 97

algorithms
 Fast Path buffer allocation algorithm 287
 Fast Path buffer allocation algorithm for BMPs 291
 Fast Path buffer allocation algorithm for CCTL threads 292
 first fit
 assigning VSO DEDB areas to data spaces 148
 HD space search algorithm 107
allocation
 IMS data sets 322
 OSAM data sets 322
alternate PCB statement 306
AM status code 475, 483
anchor point area 97
application control block (ACB)
 description 308
 generation 308
Application Control Block Maintenance utility 457
 building the ACBs 308
 deleting a data capture exit routine online 459
application development
 IMS Application Development Facility II 313
 VisualAge Generator 313
application programs
 basic load program 331
application requirements, analyzing 4, 49, 57
AREA statement 296
areas
 adding online 463
 copying data sets 119
 DEDB
 design guidelines 272
 opening 115
 preopening 115
 reopening 115
 starting 116
 stopping 116
 deleting online 463
 disabling preopen process 116
 emergency restart
 reopening 115
 FPOP= 115
 introduction to DEDB areas 114
 reopening
 emergency restart 115
 replicating data sets 119
 UOW structural definition 460
 VSO DEDB
 defining 139
Asynchronous Data Capture
 description 18
 procedure for adding 453
 using 453
auxiliary storage requirements for MSDBs 283
available length (AL) field 97

B

- background write 264
- backspacing 68
- basic initial load program, writing 329
- Basic Sequential Access Method
 - See See BSAM (Basic Sequential Access Method)
- batch message processing (BMP) programs
 - Fast Path buffer allocation algorithm 291
- batch message processing programs (BMPs)
 - accessing DEDBs 424
 - and CCTL threads 291
 - data sharing 60
 - DBCTL environment 60
 - normal buffer allocation 290
 - OBA values 423
 - overflow buffer allocation 291
 - updates in a sync interval 425
- BGWRT parameter 264
- bidirectional physically paired logical relationship 159
- bidirectional virtually paired logical relationship 161
- bit map block
 - HALDB partitions 96
- bit maps
 - calculating space 321
 - description 96
- bits in delete byte 469
- block-level data sharing 111
 - CI reclaim 241
 - SHISAM restriction 241
- blocks
 - calculating number needed 318
 - determining size 66
 - determining size of 252
 - HIDAM (Hierarchical Indexed Direct Access Method) 101
 - HISAM (Hierarchical Indexed Sequential Access Method) 70
 - PHIDAM 101
- BMP (batch message processing) programs
 - Fast Path buffer allocation algorithm 291
- BMPs (batch message processing programs)
 - accessing DEDBs 424
 - and CCTL threads 291
 - data sharing 60
 - DBCTL environment 60
 - normal buffer allocation 290
 - OBA values 423
 - overflow buffer allocation 291
 - updates in a sync interval 425
- BSAM (Basic Sequential Access Method)
 - access to GSAM databases 80
 - access to HSAM databases 65
 - access to OSAM databases 513
 - access to SHSAM databases 79
- BSIZ parameter 287, 290
- buffer handler 253
- buffer pools
 - description 253
 - designing a Fast Path 286
 - Fast Path, use 423
 - in DBCTL environment 290

- buffer pools (*continued*)
 - lookaside option 150
 - private
 - description 144
 - size determination for Fast Path 288
 - size for Fast Path determination 292
- buffers
 - allocation in Fast Path 294
 - choosing options 253
 - description 278
 - description of 257
 - Fast Path buffer allocation algorithm 287
 - Fast Path buffer allocation algorithm for BMPs 291
 - Fast Path buffer allocation algorithm for CCTL threads 292
 - fixing in storage 256, 266
 - Hiperspace buffering for VSAM 254
 - OSAM buffer sizes 256
 - specifying 256
 - system buffer allocation 288, 292
 - VSAM buffer sizes 255
- BWO(TYPEIMS) 267
 - KSDS 267
- bytes operand 98
- BYTES parameter 180, 203

C

- cache structures
 - defining a cache structure name for VSO
 - DEDB 144
 - registering name with DBRC 145
 - VSO DEDB areas 139
- calculating space
 - See space calculations
- calls
 - See also DL/I calls
 - CHKP
 - benefits in GSAM databases 80
 - benefits in SHISAM databases 80
 - UOW size considerations 274
 - GU or GN
 - See DL/I calls
 - ROLB 288, 292
 - SYNC 274
- CCTL
 - threads
 - Fast Path buffer allocation algorithm 292
- CCTL threads
 - Fast Path buffer allocation algorithm 292
- CFRM (coupling facility resource management)
 - estimating CFRM list structure size 154
- CFRM policy for MADSIOT 154
- CFSizer
 - estimating CFRM list structure size 154
- CFSTR1/2 naming convention 144
- changing a randomizing routine 457
- child segment, definition 7
- CHKP call
 - benefits in GSAM databases 80
 - benefits in SHISAM databases 80

- CHKP call (*continued*)
 - UOW size considerations 274
- CI (control interval)
 - calculating number needed 318
 - CI reclaim for VSAM KSDS 241
 - contention 424
 - DEDB (data entry database) 123
 - determining size of 252
 - HIDAM (Hierarchical Indexed Direct Access Method) 101
 - HISAM (Hierarchical Indexed Sequential Access Method) 70
 - number 99
 - overhead 317
 - PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) 101
 - SDEP 274
 - size determination in DEDB 273
 - size, changing 464
 - splits 73
- CI reclaim
 - block-level data sharing 241
 - deleting records 241
 - introduction 241
 - KSDS reorganization 241
 - VSAM REPRO, using 241
 - XRF environments 241
- CICS (Customer Information Control System)
 - background write 265
 - database types not supported 12, 60
 - DL/I Test Program 314
 - overview of access to IMS 3
 - security 31
 - sequential buffering
 - benefits 258
 - SB Initialization exit routine 264
 - using 262, 263
 - virtual storage 261
 - VisualAge Generator 313
 - VSAM database buffers 266
- CICS-DBCTL
 - GSAM 82
 - SHISAM 82
 - SHSAM 82
- CIDF (control interval definition field) 318
- code inspections 28
- coexistence considerations for HALDB Online Reorganization 375
- commands
 - allocating VSAM data sets with DEFINE CLUSTER 322
 - GENJCL.CA for HALDB OLR 387
 - GENJCL.RECOV for HALDB OLR 387
 - modifying and tuning HALDB Online Reorganization 380
 - monitoring HALDB Online Reorganization 380
 - specifying optional VSAM functions with DEFINE CLUSTER 267
 - specifying VSAM index options with DEFINE CLUSTER 269
 - starting areas with /START AREA 116
- commands (*continued*)
 - starting DEDBs with /START DATABASE 116
 - starting HALDB Online Reorganization 379
 - stopping areas with /STOP AREA 116
 - stopping DEDBs with /STOP DATABASE 116
- common synchronization point process, 425
- compressing segment data 218
- compression facility
 - See See segment edit/compression exit routine
- COMPRTN parameter
 - DBD SEGM statement 458
- concatenated key
 - converting 454
 - fields 201
 - in symbolic pointing 195
 - logical parent's 163
- concatenated segments 168, 177
- constant field 200
- CONSTANT parameter 211
- control interval (CI)
 - calculating number needed 318
 - CI reclaim for VSAM KSDS 241
 - contention 424
 - DEDB (data entry database) 123
 - determining size of 252
 - HIDAM (Hierarchical Indexed Direct Access Method) 101
 - HISAM (Hierarchical Indexed Sequential Access Method) 70
 - number 99
 - overhead 317
 - PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) 101
 - SDEP 274
 - size determination in DEDB 273
 - size, changing 464
 - splits 73
- control interval definition field (CIDF) 318
- control interval update sequence number (CUSN) 123
- conventions
 - naming
 - general rules 21
 - HALDB (High Availability Large Database) 22
 - HALDB data sets 23
- copying phase of HALDB Online Reorganization 372
- counter
 - in logical relationships 170
- counter area
 - of segment, introduction 15
- coupling facility
 - cache structure 139
 - MADSIOT 154
 - structures 144
 - structures, naming convention 144
- CP (free space chain pointer) field 97
- Cross System Product
 - See See VisualAge Generator
- crossing a logical relationship 183, 186
- cursor
 - cursor-active status for HALDB Online Reorganization 371

- cursor (*continued*)
 - HALDB Online Reorganization 373
- CUSN (control interval update sequence number) 123
- Customer Information Control System (CICS)
 - overview of access to IMS 3

D

- DA status code 475, 483
- DASD
 - contention in Fast Path 422
 - out-of-space for DEDB 425
 - space release 484
- data
 - XML
 - overview of storing in IMS databases 244
- Data Capture exit routine 459
 - adding 459
 - and logical databases 225
 - call functions 224
 - call sequence 223
 - changing 459
 - data capture exit routine 224
 - deleting online 459
 - description 17, 221
 - function 221
 - specifying in DBD 221
 - using 221, 453
- Data Dictionary
 - See DB/DC Data Dictionary 18
- data elements in segment 15
- data entry database (DEDB)
 - record distribution
 - impact of changes to UOW structure 463
- data entry databases (DEDBs)
 - adding 461
 - and segment edit/compression exit routine 218
 - areas
 - adding online 463
 - deleting online 463
 - areas, copying 119
 - areas, overview 114
 - buffer pools 290
 - CI resource contention 424
 - copying
 - areas 119
 - data sharing 119
 - DBCTL support 60
 - DEDB physical format 114
 - deleting 461
 - designing 271
 - DL/I calls against 131
 - free space algorithm 130
 - HSSP processing of 283
 - insert algorithm 129
 - IOVF
 - extending online 464
 - loading the database 338
 - modifying
 - adding and deleting segments 462
 - overview 113

- data entry databases (DEDBs) (*continued*)
 - performance considerations 422
 - replicating
 - areas 119
 - segments, adding and deleting 462
 - SSA restrictions 131
 - storage of records 126
 - when to use 113
- Data Extraction, Processing, and Restructuring system 313
- data part of segment 14, 15
- data requirements, analyzing 49, 57
- data sensitivity 190
- data set groups
 - changing the number of 418
 - multiple 18
- data set statement
 - description 296
 - HALDB (High Availability Large Database) 296
- data sets
 - allocation 322
 - copying
 - DEDB areas 119
 - DFSVSAMP 73
 - ESDS in HD databases 95
 - ESDS in secondary indexes 198
 - groups
 - changing the number of 418
 - HALDB Online Reorganization
 - naming conventions 378
 - output data sets 379, 551
 - HALDB partitions
 - maximum number of data sets 303
 - HISAM 69
 - KSDS in secondary indexes 198
 - maximum size
 - OSAM 83, 513
 - VSAM 83
 - MSDBCP1 and MSDBCP2 283
 - MSDBDUMP data set 283
 - naming conventions
 - HALDB (High Availability Large Database) 23
 - HALDB Online Reorganization 378
 - HALDB Online Reorganization overview 24
 - PHDAM 23
 - PHIDAM 23
 - PSINDEX 23
 - OSAM
 - maximum size 83, 513
 - OSAM in HD databases 95
 - pre-formatting space 267
 - replicating
 - DEDB areas 119
 - VSAM
 - maximum size 83
- data sharing
 - DEDB 119
 - VSO DEDB Areas 149
- data space
 - z/OS
 - accessing for VSO DEDB areas 148

- data space (*continued*)
 - z/OS (*continued*)
 - acquiring for VSO areas 147
- data structures, developing 49, 57
- database definition
 - HALDB partitions 299
 - using the Partition Definition utility 299
- database description (DBD)
 - coding 295
 - introduction 18
 - using dictionary to generate 18
- Database Prefix Resolution utility (DFSURG10) 357
- Database Prefix Update utility (DFSURGP0) 358
- Database Preorganization utility (DFSURPR0) 356
- database records
 - definition 6
 - HDAM (Hierarchical Direct Access Method) 98
 - HIDAM 100
 - HISAM (Hierarchical Indexed Sequential Access Method) 70
 - HSAM (Hierarchical Sequential Access Method) 65
 - introduction to 12
 - locking 109
 - MSDB (main storage database) 134
 - PHDAM (Partitioned Hierarchical Direct Access Method) 98
 - PHIDAM 100
- Database Scan utility (DFSURGS0) 356
- Database Surveyor utility (DFSPRSUR) 361
- database types
 - changing database types 394
 - changing DL/I access methods 394
 - converting database types 394
- databases
 - administration overview 3
 - application program's view 18
 - changing database types 394
 - changing DL/I access methods 394
 - CICS local-DL/I 60
 - converting database types 394
 - DBCTL support 60
 - DEDB description 113
 - defining 18
 - design
 - aids for testing 313
 - designing
 - overview 4
 - developing test databases
 - Data Extraction, Processing, and Restructuring System 313
 - DL/I test program (DFSDDLT0) 314
 - File Manager for z/OS for IMS Data 314
 - IMS Application Development Facility II 313
 - Visual Age Generator 313
 - Fast Path
 - design considerations 271
 - initial loads 329
 - full-function
 - design considerations 245
 - GSAM description 80
- databases (*continued*)
 - HALDB (High Availability Large Database)
 - description 82
 - HD description 82
 - HSAM description 64
 - implementing
 - overview 5
 - implementing database design 295
 - introduction to 11
 - loading 325
 - Fast Path initial loads 329
 - overview 5
 - restartable load program, using UCF 332
 - sample JCL 331
 - loading, description 315
 - Local-DL/I support 60
 - logical 168
 - modifying 429
 - overview 5
 - monitoring 341
 - overview 5
 - MSDB description 132
 - multiple data set groups 239
 - PCB statement 306
 - protecting during reorganization 348
 - records
 - calculating size 315
 - recovery
 - overview 5
 - reorganizing 347
 - security
 - establishing 31
 - for application programs 18
 - introduction 6
 - SHISAM description 79
 - SHSAM description 79
 - standards and procedures 6
 - testing 5, 311
 - tuning 347
 - changing the number of data set groups 418
 - overview 5
 - XML
 - overview of storing XML data 244
- DATASET statement
 - example 240
 - in logical DBD 183
- DB/DC Data Dictionary
 - establishing security 33
 - generating DBDs 18
 - generating PSBs 18
 - introduction 18
- DBBF parameter
 - DEDB Buffer Pool in the DBCTL environment 290
 - DEDB or MSDB Buffer Pools 286
- DBCTL
 - accessing from CICS, overview 3
 - CICS applications 60
 - DBBF parameter 290
 - designing DEDB buffer pools 290
- DBD (database description)
 - coding 295

- DBD (database description) *(continued)*
 - introduction 18
 - logical relationships 177
 - specifying use
 - Data Capture exit routine 221
 - field-level sensitivity 227
 - logical relationships 178, 181, 182, 183
 - multiple data set groups 239
 - secondary indexes 211
 - segment edit/compression exit routine 220
 - variable-length segments 215
 - using dictionary to generate 18
- DBD statement 181, 296
- DBDGEN (Database Description Generation)
 - utility 295
- DBDLIB 295
 - member 521
 - selecting for HALDB (High Availability Large Database) 545
- DBFDBMA0 (MSDB Maintenance utility) 133
- DBFUHDR0 (High-Speed DEDB Direct Reorganization utility) 274
- DBFX parameter 287, 290
- DBFX value 289, 293
- DCCTL
 - data sharing 60
 - GSAM (Generalized Sequential Access Method) 82
 - SHISAM (Simple Hierarchical Indexed Sequential Access Method) 82
 - SHSAM (Simple Hierarchical Sequential Access Method) 82
- DD name
 - naming conventions
 - HALDB Online Reorganization overview 23
- DDATA parameter 203
- DDNAME
 - HALDB (High Availability Large Database) 303, 520
- deactivation, record 118
- decomposed storage of XML data
 - overview 244
- DEDB (data entry database)
 - areas
 - design guidelines 272
 - functions 114
 - record distribution
 - impact of changes to UOW structure 463
- DEDB Area Data Set Create utility (DBFUMRIO)
 - copying area data sets 119
- DEDB areas
 - disabling preopen process 116
 - emergency restart
 - reopening 115
 - FPOPN= 115
 - opening 115
 - preopen
 - concurrent to operation 115
 - preopening 115
 - reopening
 - emergency restart 115
 - restarting
 - after IRLM failure 116
- DEDB areas *(continued)*
 - starting 116
 - stopping 116
 - UOW structural definition, changing 460
- DEDB CI resource
 - and DBFX value 290, 294
 - determine resource size 252
 - Fast Path Performance 421
 - overhead needed 317
- DEDB segments
 - segment growth 220
- DEDBs (data entry databases)
 - adding using online change 461
 - and segment edit/compression exit routine 218
- areas
 - adding online 463
 - deleting online 463
- areas, overview 114
- buffer pools 290
- CI resource contention 424
- data sharing 119
- DBCTL support 60
- DEDB physical format 114
- deleting using online change 461
- designing 271
- DL/I calls against 131
- free space algorithm 130
- HSSP processing of 283
- insert algorithm 129
- IOVF
 - extending online 464
- loading the database 338
- modifying
 - adding and deleting segments 462
- overview 113
- performance considerations 422
- segments, adding and deleting 462
- SSA restrictions 131
- storage of records 126
- when to use 113

- DEFINE CLUSTER command
- for VSAM index option 269
- in access method services 267
- VSAM data set allocation 322
- defining data set groups
- HALDB (High Availability Large Database) 303
- delete byte
- bits 469
- description 15
- HDAM 100
- HISAM 70
- HSAM 67
- in logical relationships 483
- in secondary indexes 200
- PHDAM (Partitioned Hierarchical Direct Access Method) 100
- delete rules for logical relationships 188, 189, 481, 511
- deleting segments
- HD databases 107
- HISAM databases 76
- HSAM databases 68

- dependent segment, definition 7
- design aids
 - for test databases 313
- design reviews
 - description of 25
 - introduction 4
- destination parent 169, 190
- determining VSAM options 264
- DFXPMB 546
- DFXPMB Partition Definition utility panel 545
- DFSCTL data set control statements
 - SB control statement 262
 - SBPARM control statement 262
- DFSDDLTO (DL/I test program) 314
- DFSMTB0 (DB Monitor program) 341
- DFSPRCT1 (Partial Database Reorganization utility) 362
- DFSPRSUR (Database Surveyor utility) 361
- DFSUOCU0 (Online Change Copy utility)
 - deleting a data capture exit routine 459
- DFSUOCU0 (Online Change utility) 459
 - changing a randomizing routine 457
- DFSURG10 (Database Prefix Resolution utility) 357
- DFSURGL0 (HD Reorganization Reload utility) 355
- DFSURGP0 (Database Prefix Update utility) 358
- DFSURGS0 (Database Scan utility) 356
- DFSURGU0 (HD Reorganization Unload utility) 354
- DFSURPR0 (Database Preorganization utility) 356
- DFSURRL0 (HISAM Reorganization Reload utility) 354
- DFSURUL0 (HISAM Reorganization Unload utility) 353
- DFSVSAMP data set 73
- DFSVSMMxx member of IMS.PROCLIB
 - MADSIOT 154
- dictionary
 - See DB/DC Data Dictionary
- direct access methods
 - HDAM (Hierarchical Direct Access Method) 82
 - HIDAM (Hierarchical Indexed Direct Access Method) 82
 - PHDAM (Partitioned Hierarchical Direct Access Method) 82
 - PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) 82
- direct address pointers 82, 85
- direct dependent segment types (DDEP) 126
- direct pointers
 - logical relationships 162, 164, 167, 189
 - secondary indexes 200, 201
- direct storage method 60
- distribution of database records
 - DEDB 463
- DL/I
 - access methods
 - changing 394
 - changing from HDAM to PHDAM and HIDAM to PHIDAM 401
 - converting 394
 - from PHDAM and PHIDAM to HDAM and HIDAM 402
 - HDAM to HIDAM 399
 - HDAM to HISAM 398

- DL/I (continued)
 - access methods (continued)
 - HIDAM to HDAM 397
 - HIDAM to HISAM 397
 - HISAM to HDAM 395
 - HISAM to HIDAM 395
 - and ACBs 308
 - calls
 - DEDBs 131
 - HD databases 84
 - HISAM databases 72
 - HSAM databases 67
 - MSDB 135
 - calls in logical relationships
 - delete call 483
 - logical child insert call 472
 - replace call 475
 - DL/I Call Summary report 408
- DL/I access methods
 - changing 394
 - from HDAM to PHDAM and HIDAM to PHIDAM 401
 - from PHDAM and PHIDAM to HDAM and HIDAM 402
 - HDAM to HIDAM 399
 - HDAM to HISAM 398
 - HIDAM to HDAM 397
 - HIDAM to HISAM 397
 - HISAM to HDAM 395
 - HISAM to HIDAM 395
 - converting 394
- DL/I Call Summary report 408
- DL/I calls
 - DEDBs 131
 - HD databases 84
 - HISAM databases 72
 - HSAM databases 67
 - in logical relationships
 - delete call 483
 - logical child insert call 472
 - replace call 475
 - MSDB 135, 138
- DL/I test program (DFSDDLTO) 314
- DLET call
 - DASD space release 484
- DLIModel utility
 - storing XML data
 - overview 244
- DREF (disabled reference) option
 - for VSO-area data spaces 147
- DSPXPDA 543
- DSPXPEA 543
 - database name 543
 - output data set name 544
- DSPXPPIA 544
 - database name 544
 - input data set name 544
 - input member name 544
 - processing option 544
- DSPXPKE panel 532
- DSPXPLA 535

DSPXPLA (continued)

- act 535
- data set name prefix 535
- ID 535
- name 535
- DSPXPLB 536
- DSPXPOA 542
- DSPXRUN command 548
 - database_name 548
 - dataset_name 548
 - member_name 548
 - processing_option 548
- dump option 266
- DUMP parameter 266, 269
- duplex paths 482
- duplicate data field 201
- duplicate data in logical relationships 157
- duplicate keys 198
- DX status code 483

E

- ECNT (extended communications node table) 136
- edit/compression exit routine
 - See segment edit/compression exit routine
- editing segment data 218
- emergency restart
 - DEDB areas
 - reopening 115
- encoding data
 - See segment edit/compression exit routine
- encrypting data 33
- END statement 298, 307
- Error Queue Element (EQE) 117
- ESAF
 - See external subsystem attach facility
- ESCD (extended system contents directory) 136
- ESDS (entry-sequenced data set)
 - HD databases 95
 - HISAM 69
 - secondary indexes 198
- estimating minimum database size 252
- EXIT parameter 222
- exit routines
 - changing, online 457
- exporting database definitions
 - HALDB (High Availability Large Database) 303
- extended communications node table (ECNT) 136
- extended system contents directory (ESCD) 136
- external subsystem attach facility 61
- EXTRTN parameter 204, 212

F

- fallback considerations for HALDB Online
 - Reorganization 376
- Fast Path
 - access to DL/I databases 131
 - buffers 423
 - CI contention 345, 424
 - committing updates 153

Fast Path (continued)

- common sync point processing 426
- control interval 424
- data sharing
 - DEDB 119
- databases
 - DEDB overview 113
 - MSDB overview 132
 - overview 113
- environments 113
- Fast Path buffer allocation algorithm 287
- Fast Path buffer allocation algorithm for BMPs 291
- Fast Path buffer allocation algorithm for CCTL
 - threads 292
- initial database load 329
- interpreting analysis reports 345
- loading the database 337
- log analysis 343
- log reduction 344
- mixed mode 131
- monitored events 345
- monitoring and tuning 343
- output thread 153
- performance considerations 343
- Resource Name Hash routine 427
- selecting transactions 345
- subset pointers 127, 277
- synchronization point processing 153, 425
- transaction timings 344
- tuning Fast Path systems 421
- user hash routine, programming considerations 427
- using the Log Analysis utility (DBFULTA0) 343

- Fast Path virtual storage option
 - See virtual storage option
- fbff (free block frequency factor) 245
- FCP (forward chain pointer) 134
- FH status code 117
- FID (fixed intersection data) 171
- FIELD statement
 - definition 202
 - in secondary indexing 213
 - in the DBD 269
 - position in DBD 297
- field-level sensitivity
 - description of 226
 - establishing security 31
 - inserting segments 229
 - introduction 17
 - overlapping paths 230
 - path calls 230
 - replacing segments 228
 - retrieving segments 228
 - specifying in DBD and PSB 227
 - use with variable-length segments 231
 - uses 226
 - using 226
- fields 201
 - AL 97
 - constant 200
 - CP 97
 - definition 6

- fields *(continued)*
 - duplicate data 201
 - FSE 97
 - FSEAP 96
 - ID 97
 - in segment 15
 - pointer 200
 - search 200
 - subsequence 200
 - system related 202
- File Action bar 520
- FINISH statement 298
- first fit algorithm to assign VSO DEDB areas to data spaces 148
- fixed intersection data (FID) 171
- fixed-length segments
 - specifying minimum size 220
- fixed-length segments, definition 14
- FLD (Field) call 138
- format
 - CI in DEDB 123
 - DEDB segments 123
 - fixed-length segments 14
 - HD databases 95
 - HDAM segments 100
 - HIDAM index segment 102
 - HIDAM segments 101
 - HISAM segments 70
 - HSAM segments 66
 - PHDAM segments 100
 - PHIDAM index segment 102
 - PHIDAM segments 101
 - pointer segment 199
 - variable-length segments 14
- formula
 - estimating CFRM list structure size 154
 - first fit algorithm 148
- formulas for
 - calculating buffers for Fast Path 288, 292
 - calculating space for MSDBs 283
 - calculating storage for MSDB 278
 - size of root addressable area 246
- forward chain pointer 134
- FPOPN=
 - overview 115
- FPRLM=
 - restarting DEDB areas 116
- FR status code
 - for BMP regions 289
 - for CCTL threads 293
 - in Fast Path buffer allocation 288
 - in Fast Path buffer allocation for BMPs 292
- free block frequency factor (fbff) 245
- free logical record 72
- free space
 - chain pointer (CP) field 97
 - element (FSE) 97
 - element anchor point (FSEAP) 96
 - HD (Hierarchical Direct) 96
 - HDAM (Hierarchical Direct Access Method) 245
 - HIDAM 245

- free space *(continued)*
 - HIDAM (Hierarchical Indexed Direct Access Method) 101
 - KSDS 267
 - percentage factor 246
 - PHDAM (Partitioned Hierarchical Direct Access Method) 245
 - PHIDAM 245
 - PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) 101
 - space calculations 321
- FREESPACE parameter 267
- FRSPC parameter 245
- FS status code 275
- FSE (free space element) 97
- FSEAP (free space element anchor point) 96
- fspf (free space percentage factor) 246
- full-duplex paths 482
- full-function segments
 - specifying minimum size 220
- FW status code
 - for CCTL threads 293
 - in BMP regions 289
 - in Fast Path buffer allocation 288
 - in Fast Path buffer allocation for BMPs 292

G

- GC status code 274, 285
- GE status code 177
- general format of HD databases and use of special fields 321
- Generalized Sequential Access Method (GSAM)
 - See GSAM (Generalized Sequential Access Method) 78
- GPSB (Generated PSB)
 - I/O PCB 309
 - modifiable alternate response PCB 309
- GSAM (Generalized Sequential Access Method) 78, 80, 337

H

- HALDB (High Availability Large Database) 82
 - adding partitions 302
 - automatic partition definition 302
 - automatic partition definition using Partition Definition utility 527
 - batch import 303
 - bit map block for partition 96
 - Change Partition screen 530
 - F11 532
 - F5 532
 - F6 532
 - Partition high key 531
 - Partition ID field 531
 - Partition Selection String 531
 - changing 404
 - HALDB Partition Selection exit routine 405
 - overview 404
 - partition boundaries 406

HALDB (High Availability Large Database) *(continued)*

- changing *(continued)*
 - partition key ranges 406
 - partition structure modification 405
 - single partitions 404
- changing DL/I access methods
 - changing from HDAM to PHDAM and HIDAM to PHIDAM 401
 - from PHDAM and PHIDAM to HDAM and HIDAM 402
- changing partitions using the PDU 300
- configuration
 - list 546
- copying partitions 302
- creating HALDB (High Availability Large Database) partitions 299
- creating with the Partition Definition utility (PDU) 299
- data set naming conventions 23, 378
- data set statement 296
- data sets
 - maximum per partition 303
- Database Partition list
 - act 537
 - anchor 537
 - bytes 537
 - FBFF 537
 - File action bar choice 535
 - FSFF 538
 - high block 537
 - module 537
 - name 537
- Database Partitions list 534
 - displaying 534
- DBDLIB 304
- defining data set groups 303
- defining data set groups using Partition Definition utility 533
- defining with the Partition Definition utility (PDU) 299
- definition process 299
- Delete a Database panel using Partition Definition utility 543
- deleting a database 304
- deleting partitions 302
- DSPXPDA 543
- DSPXPEA
 - See DSPXPEA 543
- DSPXPGA 533
- DSPXPGB 534
- DSPXPGC 534
- DSPXPJA
 - See DSPXPJA 544
- DSPXPKE panel 532
- DSPXPLA
 - See DSPXPLA 534
- DSPXPLB 536
- DSPXPOA 542
- DSPXRUN command 548
- exporting database definitions 303

HALDB (High Availability Large Database) *(continued)*

- fallback
 - to HDAM and HIDAM 402
- File Action bar 520
 - actions 521
- finding partitions 302
- HALDB (High Availability Large Database),
 - defining 299
- HALDB Online Reorganization 370
- HALDB Partition Definition utility (%DFSHALDB)
 - accessing Help 519
 - exiting the utility 519
 - main panel 518
 - main screen 518
 - Partitioned Databases panel 518
 - using 517
- HALDB Partition Definition utility (PDU) 299
- HALDB Partition Selection exit routine (DFSPSE00)
 - changing 405
 - modifying 405
 - replacing 405
- ILDS, updating
 - offline reorganization 369
- importing database definitions 303
- importing database information using Partition Definition utility 544
- IMS concatenation
 - displaying using Partition Definition utility 544
- IMS configuration 545
- indirect list data set (ILDS)
 - allocating 304
- indirect list entry (ILE)
 - description 304
- indirect list key (ILK)
 - description 304
- information
 - changing using Partition Definition utility 521
 - deleting using Partition Definition utility 521, 543
 - exporting using Partition Definition utility 521, 543
 - importing using Partition Definition utility 521, 544
 - opening using Partition Definition utility 542
 - viewing using Partition Definition utility 521
- interfaces
 - HALDB Partition Definition utility (%DFSHALDB) 517
- ISRDDN 544
- LCHILD statement 298
- line commands 538
- manual partition definition 302
- manual partition definition using Partition Definition utility 528
- master
 - values 521
- maximum size 83
- migrating
 - fallback to HDAM and HIDAM 402
 - from HDAM to PHDAM and HIDAM to PHIDAM 401
- modifying 404

HALDB (High Availability Large Database) *(continued)*

- HALDB Partition Selection exit routine 405
- overview 404
- partition boundaries 406
- partition key ranges 406
- partition structure modification 405
- single partitions 404
- modifying data set groups 303
- modifying partitions 302
- naming conventions 22
- offline reorganization 365
 - overview 365
 - reallocating data sets 368
 - reloading partitions 369
 - unloading partitions 367
 - updating ILDS 369
- online reorganization 370
 - DD name naming convention overview 23
 - naming convention 378
 - naming convention overview 24
- overview 82
- partition bit map block 96
- partition definition 299
- Partition Definition utility 299, 529
 - accessing 517
 - high key value 528
 - impact on RECON 517
 - modifying fields 529
 - panels 517
 - selecting the DBDLIB and RECON configuration 545
- Partition Definition utility (PDU) 298
 - Change Partition panel 300
- partition high key 301
- partition structures, updating 405
- partitions
 - changing 404
 - changing boundaries 406
 - changing key ranges 406
 - changing using Partition Definition utility 521
 - changing, overview 404
 - copying using Partition Definition utility 538
 - creating using Partition Definition utility 521
 - data sets, maximum 303
 - defining using Partition Definition utility 521
 - deleting using Partition Definition utility 538
 - manual definition using Partition Definition utility 528
 - maximum number 521
 - modifying 404
 - modifying key ranges 406
 - modifying, overview 404
 - naming conventions 23
 - opening using Partition Definition utility 521, 538
 - printing information using Partition Definition utility 538
- pointers
 - self-healing pointer process 388
- printing partitions 302
- reallocating data sets
 - offline reorganization 368

HALDB (High Availability Large Database) *(continued)*

- RECON 304
 - RECON data set
 - selecting 545
 - reloading partitions
 - offline reorganization 369
 - reorganizing 364
 - offline 365
 - reallocating data sets 368
 - reloading partitions 369
 - secondary indexes 370
 - unloading partitions 367
 - updating ILDS 369
 - REUSE parameter 322
 - secondary indexes
 - reorganizing 370
 - self-healing pointer process 388
 - performance 392
 - sorting partitions 302
 - unloading partitions
 - offline reorganization 367
 - viewing DDNAME 303
 - viewing partitions 302
- ## HALDB (High Availability Large Database) partition definition utility
- registering OLR capability with DBRC 523
- ## HALDB Online Reorganization
- coexistence considerations 375
 - copying phase 372
 - cursor 373
 - cursor-active status 371
 - Database Change Accumulation utility 387
 - DD name naming convention
 - overview 23
 - dynamic PSB 372
 - fallback considerations 376
 - FDBR 384
 - GENJCL.CA command 387
 - GENJCL.RECOV command 387
 - image copy utilities 387
 - initialization phase 371
 - locking 384
 - log impact 383
 - migration considerations 375
 - modifying 380
 - monitoring 380
 - naming conventions
 - overview 24
 - output data set requirements 379, 551
 - overview 370
 - RATE parameter of INITIATE OLREORG command 383
 - recovery 386
 - Remote Site Recovery (RSR) 384
 - requirements for output data sets 379, 551
 - restart 383, 384
 - restrictions 376
 - sequential buffering 388
 - starting 379
 - system impact 383
 - termination phase 374

- HALDB Online Reorganization (*continued*)
 - tuning 380
 - unit of reorganization 373
 - utilities 385
 - XRF 383
- HALDB Partition Definition utility (%DFSHALDB)
 - accessing Help 519
 - exiting the utility 519
 - main panel 518
 - options 518
 - main screen 518
 - options 518
 - Partitioned Databases panel 518
 - options 518
 - using 517
- HALDB Partition Selection exit routine (DFSPSE00)
 - changing 405
 - modifying 405
 - replacing 405
- half-duplex paths 482
- HB (hierarchic backward) pointers 87
- HD Reorganization Reload utility
 - ILDS
 - control statement specifications 369
 - updating 369
- HD Reorganization Reload utility (DFSURGL0) 355
- HD Reorganization Unload utility (DFSURGU0) 354
- HD space search algorithm 107
- HD Tuning Aid 247
- HDAM
 - adjusting 411
 - options, adjusting 411
- HDAM (Hierarchical Direct Access Method)
 - accessing segments 103
 - calls against 84
 - changing DL/I access methods
 - from HDAM 397
 - from HISAM 395
 - from PHDAM 402
 - to HDAM 399
 - to HISAM 398
 - to PHDAM 401
 - database records 100
 - database records, locking 109
 - deleting segments 107
 - description of 82
 - format of database 95
 - inserting segments 104
 - loading the database 337
 - locking 111
 - logical record length 252
 - maximum size 83
 - multiple data set groups 237
 - options available 84
 - OSAM (overflow sequential access methods)
 - used 95
 - overflow area 98
 - pointers in 85
 - randomizing module 247
 - root addressable area 98, 100
 - segment format 100
- HDAM (Hierarchical Direct Access Method) (*continued*)
 - size of root addressable area 246
 - space calculations 315
 - specifying free space 245
 - storage of records 98
 - when to use 84
 - z/OS access methods used 83
- HF (hierarchic forward) pointers
 - description 86
- HIDAM (Hierarchical Indexed Direct Access Method)
 - accessing segments 103
 - calls against 84
 - changing DL/I access methods
 - from HDAM 399
 - from HISAM 395
 - from PHIDAM 402
 - to HDAM 397
 - to HISAM 397
 - to PHIDAM 401
 - deleting segments 107
 - description of 82
 - format of database 95
 - index database 83, 100
 - index segment 102
 - inserting segments 104
 - loading the database 337
 - locking 111
 - logical record length 252
 - maximum size 83
 - multiple data set groups 237
 - options available 84
 - pointers in 85
 - RAPs, using 102
 - segment format 101
 - sequential root processing 103
 - space calculations 109, 315
 - specifying free space 245
 - storage of records 100
 - when to use 85
- hierarchic
 - backward pointers 87
- hierarchic forward (HF) pointers
 - description 86
- Hierarchical Direct Access Method
 - See HDAM (Hierarchical Direct Access Method) 322
- Hierarchical Indexed Direct Access Method
 - See HIDAM (Hierarchical Indexed Direct Access Method) 322
- Hierarchical Indexed Sequential Access Method
 - See HISAM (Hierarchical Indexed Sequential Access Method) 322
- Hierarchical Indexed Sequential Access Method (HISAM)
 - accessing segments 72
 - calls against 72
- Hierarchical Sequential Access Method
 - See HSAM (Hierarchical Sequential Access Method) 322
- Hierarchical Sequential Access Method (HSAM)
 - accessing segments 67

- Hierarchical Sequential Access Method (HSAM)
 - (*continued*)
 - calls against 67
- hierarchy
 - concept explained 8
 - definition 7
 - restructuring of with secondary indexes 197
- high key
 - of HALDB partitions 301
 - value, entering 301
- High-Speed DEDB Direct Reorganization utility (DBFUHDR0) 274
- high-speed sequential processing (HSSP)
 - description 283
- hiperspace buffering 412
- HISAM (Hierarchical Indexed Sequential Access Method)
 - access method 69
 - accessing segments 72
 - calls against 72
 - changing DL/I access methods
 - from HDAM 398
 - from HIDAM 397
 - to HDAM 395
 - to HIDAM 395
 - database reorganization procedures 364
 - deleting segments 76
 - description of 69
 - inserting segments 72
 - loading the database 336
 - locking 110
 - logical record format 71
 - logical record length 249, 252
 - options available 69
 - performance 74, 78
 - pointers 71
 - replacing segments 78
 - segment format 70
 - space calculations 315
 - storage of records 69
 - when to use 69, 78
- HISAM Reorganization Reload utility (DFSURRL0) 354
- HISAM Reorganization Unload utility (DFSURUL0) 353
- HSAM (Hierarchical Sequential Access Method)
 - accessing segments 67
 - calls against 67
 - deleting segments 68
 - description of 64
 - inserting segments 68
 - options available 65
 - performance 68
 - replacing segments 68
 - segment format 66
 - space calculations 315
 - storage of records 65
 - when to use 65
 - z/OS access methods used 65
- HSSP (high-speed sequential processing)
 - description 283
 - for database recovery 286
 - image-copy option 285

- HSSP (high-speed sequential processing) (*continued*)
 - limits and restrictions 284
 - private buffer pools 286
 - processing option H 285
 - reasons for choosing 284
 - SETO statement 285
 - SETR statement 285
 - UOW locking 286
 - using 285
- I**
- I/O errors
 - ADS 154
 - MADS 154
 - multiple area data sets (MADS) 154
 - single area data sets (ADS) 154
- I/O PCB 309
- ID (task ID) field 97
- IDP and Fast Path 343
- IEFBR14 utility 322
- IEHPROGM program 322
- IFP and MPP regions
 - maintaining continuous availability of 455
- ILDS
 - reorganization updates 369
- ILDS (indirect list data set)
 - allocating 304
 - calculating size 305
 - defining 304
 - sample JCL 304
 - size, calculating 305
- ILE (indirect list entry) 304
- ILK (indirect list key) 304
- image-copy option 285
- IMBED I NOIMBED parameter 268
- implementing database design 5, 295
- importing database definitions
 - HALDB (High Availability Large Database) 303
- IMS Data Capture exit
 - See Data Capture exit routine
- IMS High Performance Pointer Checker 247
- IMS trace parameters
 - about 266
- IMS.ACBLIB library 308
 - online change procedure 457, 459
- IMS.DBDLIB 295
- IMS.PSBLIB 305
- in physical databases 182
- in the physical DBD 181
- independent overflow part of area (IOVF)
 - description 123
 - extending online 464
- index maintenance exit routine 204
- index segment 102
- index set records 268
- indexed databases 83
 - HIDAM 100
 - HISAM 69
 - PHIDAM 100
- INDICES parameter 207

- indirect list data set (ILDS)
 - allocating 304
 - calculating size 305
 - defining 304
 - sample JCL 304
 - size, calculating 305
- indirect list entry (ILE) 304
- indirect list key (ILK) 304
- initial load program
 - basic 331
 - Fast Path 329
 - restartable, using UCF 332
 - writing 329
- initialization phase of HALDB Online
 - Reorganization 371
- input for DBDGEN utility
 - DBD 295
- INSERT parameter
 - free space for a KSDS 265, 267
 - using in splitting CIs 73
- insert rules for logical relationships 188, 189, 471, 475
- insert strategy
 - choosing 265
- inserting segments
 - DEDB SDEPs 275
 - HD databases 104
 - HISAM databases 72
 - HSAM databases 68
 - MSDB (main storage database) 136
- inspections
 - code inspections 28
 - security inspection 28
- intact storage of XML data
 - overview 244
- intersection data 170, 172
- IOB (input/output block) 266
- IOBF parameter 256
- IOVF
 - See independent overflow part of Area
- IOVF (independent overflow part of area)
 - extending online 464
- IRLM
 - failure
 - restarting DEDB areas 116
- IRLM (internal resource lock manager)
 - block-level data sharing 111
 - failure
 - restarting DEDB areas 116
 - locking protocols 109
- ISPF
 - batch job 547
- ISPF member list 520
 - display 520
- ISPF panels
 - HALDB Partition Definition utility
 - (%DFSHALDB) 518
- ISPSTART 548
- ISPTLIB 543
- ISRDDN command 544
- ISRT (insert), loading a database 325
- IWAITS/CALL field 408

J

- JCL (Job Control Language)
 - for allocating data sets 322
 - for initial load program 336
- Job Control Language
 - See JCL (Job Control Language) 322

K

- KEY sensitivity 190
- key sequenced data sets (KSDS)
 - CI reclaim 241
- keys
 - ascending sequence 65
 - concatenated 201
 - duplicate 198
 - unique in secondary indexes 202
- KSDS (key sequenced data sets)
 - CI reclaim 241
- KSDS (key-sequenced data set)
 - HISAM (Hierarchical Indexed Sequential Access Method) 69
 - secondary indexes 198
 - specifying BWO(TYPEIMS) 267
 - specifying free space for 267

L

- LCF (logical child first) pointer 164
- LCHILD statement
 - description 297
 - HALDB (High Availability Large Database) 298
 - in logical relationships 178, 181
 - in secondary indexing 211
- LCL (logical child last) pointer 164
- level in hierarchy 11
- levels in VSAM index 268
- LGNR 344
- libraries
 - IMS.ACBLIB 308
 - IMS.DBDLIB 295
 - IMS.PSBLIB 305
- list structure
 - defining 154
 - estimating size 154
- LKASID
 - INIT.DBDS and INIT.CHANGE parameter 141
- LOAD (load), description 325
- load program
 - example of initial load program 331
- load program, writing 325
- loading databases
 - description 315
 - introduction 5
 - MSDB (main storage database) 281
 - sample programs 331
- local views, developing a data structure 49
- locking impact of HALDB Online Reorganization 384
- locking protocols 109
- log analysis, Fast Path information 343

- log facility, Fast Path performance 422
- log impact of HALDB Online Reorganization 383
- log reduction 344
- logic
 - for initial load program 331
 - for restartable initial load program 333
- logical child first (LCF) pointer 164
- logical child in logical relationships 158, 162
- logical child last (LCL) pointer 164
- logical databases 168
- logical DBD 182, 189
- logical parent in logical relationships 158, 162
- logical parent pointer
 - See LP (logical parent) pointer 162
- logical parent's concatenated key (LPCK) 163
- logical records
 - HD (Hierarchical Direct) 95
 - HISAM 70, 249, 252
 - overhead 318
 - secondary indexes 199
- logical relationships
 - analyzing requirements 56, 57
 - and Data Capture exit routine 225
 - bidirectional physically paired 159
 - bidirectional virtually paired 161
 - comparison with secondary indexes 214
 - concatenated segments 169
 - counter 170
 - crossing 183, 186
 - delete rule restrictions 225
 - delete rules 188, 481, 511
 - description of 157, 189
 - DLET call
 - DASD space release 484
 - DLET calls 483
 - establishing 172
 - insert rules 188, 472, 475
 - intersection data 170, 172
 - ISRT call 472
 - loading databases 337
 - logical child 158, 162
 - logical parent 158, 162
 - overview 56
 - paths 168, 169
 - performance considerations 189, 192
 - physical parent 158, 162
 - pointers 162, 167
 - procedures for adding to existing databases 433
 - REPL call 475
 - replace rules 188, 475, 479
 - requirements, analyzing 56
 - restrictions on modifying 449
 - rules 511
 - rules for defining 181, 182, 183, 189
 - secondary indexes, with 209
 - sequence fields 176, 177
 - specifying in DBD 178, 181, 182, 183
 - uses 157
 - virtual logical children 177
- logical twin backward (LTB) pointer 166
- logical twin chains 191

- logical twin forward (LTF) pointer 166
- logical twin pointer 515
- long busy 154
- lookaside option
 - for buffer pools 150
- lookaside option for buffer pools, description 150
- lookaside, defining private buffer pools 146
- LP (logical parent) pointer 162
 - correcting bad pointers 515
 - definition 162
 - performance considerations 189
- LPCK (logical parent's concatenated key) 163
- LTB (logical twin backward) pointer 166
- LTERM 132
- LTF (logical twin forward) pointer 166

M

- macros
 - PCB 295
 - PSB 295
- MADSIOT (Multiple Area Data Set I/O Timing)
 - calculating list structure storage size 154
 - CFRM 154
 - coupling facility 154
 - long busy 154
- main storage database
 - See MSDB (main storage database) 337
- main storage utilization, Fast Path 425
- maintenance
 - databases, planning 269
 - secondary indexes 205
- making keys unique using system related fields 202
- many-to-many mapping 50
- mapping data aggregates 50
- maximum size
 - HALDB (High Availability Large Database) 83
 - HDAM database 83
 - HIDAM database 83
 - PHDAM database 83
 - PHIDAM database 83
- MBR parameter 183
- migrating
 - fallback
 - from HALDB 402
 - from PHDAM and PHIDAM 402
 - to HDAM and HIDAM 402
 - from HDAM to PHDAM and HIDAM to PHIDAM 401
 - to HALDB 401
- migration considerations for HALDB Online Reorganization 375
- minimum size
 - specifying for full-function segments 220
- mixed mode 131
- mixing pointers 93
- modifiable alternate response PCB 309
- modifying a database
 - description of 429
 - introduction 5
- modifying data set groups
 - HALDB (High Availability Large Database) 303

- MON parameter 342
- monitoring
 - and tuning Fast Path systems 343
 - description of 341
 - events for Fast Path 345
 - introduction
 - overview 5
 - reports 341
- movement in hierarchy 10
- MSDB (main storage database)
 - calls against 135
 - deleting segments 136
 - description of 132
 - design considerations 277, 286
 - inserting segments 136
 - loading the database 337, 429
 - MSDB Maintenance utility (DBFDBMA0) 133
 - options available 132
 - page fixing 281
 - position 137
 - restrictions on changing DBD 429
 - storage of records 134
 - when to use 131, 133
- MSDBCP1 data set 283
- MSDBCP2 data set 283
- MSDBDUMP data set 283
- multi-area structure
 - duplexing 143
- Multiple Area Data Set I/O Timing (MADSIOT)
 - calculating list structure storage size 154
 - CFRM 154
 - coupling facility 154
 - long busy 154
- multiple area data sets (MADS)
 - I/O errors 154
 - MADSIOT 154
- multiple data set groups
 - description of 235
 - HD databases 237
 - introduction 18
 - specifying in DBD 239
 - storage of records 238
 - uses 236
 - using 235

N

- NAME parameter
 - in a DBD 183, 211
 - in the SENFLD statement 227
- naming convention
 - examples of defining 144
- naming convention, coupling facility structure 144
- naming conventions 21
 - general rules 21
 - HALDB (High Availability Large Database) 22
 - HALDB data sets 23
 - HALDB online reorganization
 - DD name overview 23
 - overview 24

- NBA (normal buffer allocation)
 - for CCTL 290
 - in DBCTL environment 290
 - limit 289
 - use of 287
- NBA parameter 278
- NBA/FPB limit 293
- NBRSEGS parameter 282
- NE status code 206
- no free logical record 73
- NOLKASID
 - INIT.DBDS and INIT.CHANGE parameter 141
- non-terminal-related database 132
- NOPROT parameter 206
- normal buffer allocation (NBA)
 - for CCTL 290
 - in DBCTL environment 290
 - use of 287
- NULLVAL parameter 204, 212

O

- OBA (overflow buffer allocation)
 - for CCTL threads 291
 - in DBCTL environment 291
 - use of 287
- OLR (online reorganization)
 - HALDB (High Availability Large Database)
 - registering OLR capability with DBRC using PDU 523
- one-to-many mapping 50
- online change
 - adding a DEDB 461
 - adding DEDB areas online 463
 - databases 454
 - deleting a DEDB 461
 - deleting DEDB areas online 463
 - exit routines 457
 - randomizing routines 457
- Online Change Copy utility (DFSUOCU0)
 - deleting a data capture exit routine 459
- Online Change utility (DFSUOCU0)
 - changing a randomizing routine 457
- online reorganization
 - HALDB
 - data set naming convention overview 24
 - DD name naming convention overview 23
 - HALDB naming convention 378
 - HALDB Online Reorganization 370
 - online reorganization (OLR)
 - HALDB (High Availability Large Database)
 - registering OLR capability with DBRC using PDU 523
- opening
 - DEDB areas 115
- operands
 - /CK 202
 - /SX 202
 - See parameters
- optional functions
 - Data Capture exit routines 221

- optional functions *(continued)*
 - field-level sensitivity 226
 - GSAM databases 81
 - HD databases 84
 - HISAM databases 69
 - HSAM (Hierarchical Sequential Access Method) 65
 - logical relationships 157, 189
 - MSDB databases 132
 - multiple data set groups 235
 - secondary indexes 192
 - Segment Edit/Compression exit routine 218
 - SHISAM databases 80
 - SHSAM databases 79
 - variable-length segments 215
- OPTIONS statement
 - fixing buffers in VSAM 256
 - for OSAM 269
 - for VSAM 264, 266
 - OSAM 269
 - use in splitting CIs 73
- OSAM
 - data sets
 - maximum size 83
- OSAM (Overflow Sequential Access Method)
 - adjusting buffers 413
 - allocation of data sets 322
 - description 257, 513
 - options 269
 - track space used 252
 - used by HD 95
- OSAM data sets
 - maximum size 513
- OSAM Sequential Buffering (SB)
 - See SB (OSAM Sequential Buffering) 257
- output thread 153
- overflow buffer allocation (OBA)
 - See OBA (overflow buffer allocation) 291
- overflow data set 69
- Overflow Sequential Access Method
 - See OSAM (Overflow Sequential Access Method) 513
- overflow space allocation
 - changing 464
- overhead
 - DEDB CI resources 317
 - logical records 318

P

- packing density 248
- page fixing MSDBs 281
- parameters
 - BGWRT 264
 - BSIZ
 - in DB/TM environment 287
 - in the DBCTL environment 290
 - BWO(TYPEIMS) 267
 - BYTES 203
 - CNBA 291
 - CONSTANT 211
 - DB Monitor 342

- parameters *(continued)*
 - DBBF
 - in DB/TM environment 286
 - in the DBCTL environment 290
 - DBFX
 - in DB/TM environment 287
 - in the DBCTL environment 290
 - DDATA 203
 - DUMP 266, 269
 - EXIT 222
 - EXTRTN 204, 212
 - FPB 291
 - FPOB 291
 - FREESPACE 267
 - FRSPC 245
 - IMBED | NOIMBED 268
 - INDICES 207
 - INSERT
 - free space for a KSDS 265, 267
 - using in splitting CIs 73
 - IOBF 256
 - LGNR 344
 - MBR 183
 - MON 342
 - NAME
 - in a DBD 183, 211
 - in the SENFLD statement 227
 - NBA 278
 - NBRSEGS 282
 - NOPROT 206
 - NULLVAL 204, 212
 - PARENT 169, 183
 - in logical relationships 180, 183
 - to specify PCF and PCL pointers 90
 - to specify PCF pointers 89
 - PASSWD 33
 - POINTER 181
 - PROCOPT 32, 275
 - PROCSEQ 194, 197
 - PROT 206
 - RECORD 252
 - REPL 228
 - REPLICATE | NOREPLICATE 268
 - RMNAME 98
 - HDAM options 248
 - PHDAM options 248
 - specifying number of blocks or CIs 247
 - specifying number of RAPS 97
 - RULES 471, 511
 - SEGMENT 211
 - SHAREVL 120
 - SOURCE 181, 190
 - SPEED | RECOVERY 267
 - SRCH 212
 - START 203
 - SUBSEQ 202, 212
 - TYPE 228
 - VERSION 222
 - VSAMFIX 256, 266
 - VSAMPLS 266
- PARENT parameter 89, 169, 180, 183

- parent segment, definition 7
- Partial Database Reorganization utility (DFSPRCT1) 362
- Partition Default information screen
 - anchor 525
 - automatic definition 524, 527
 - block size 526
 - bytes 525
 - data set name prefix 525
 - database name 524
 - default JCL 526
 - free block freq. factor 526
 - free space percentage 526
 - high block number 525
 - image copy JCL 527
 - input data set 525
 - max. image copies 526
 - module name 525
 - online image copy JCL 527
 - partition ID 525
 - receive JCL 527
 - recovery period 526
 - recovery utility JCL 526
 - reusable? 527
 - use defaults for DS groups 525
- partition definition utility
 - HALDB (High Availability Large Database)
 - registering OLR capability with DBRC 523
- Partition Definition utility (PDU)
 - changing partitions 300
 - creating HALDB partitions 299
 - HALDB functions 298
 - high key value, entering 301
 - partition definition steps 299
 - partition high key value, entering 301
- partition high key 301
 - entering the high key value 301
- partition structure modification 405
- partitioned database 82
 - information screen
 - database name 522
 - database organization 522
 - number of data set groups 523
 - part. selection routine 522
 - recoverable? 523
 - RSR global service group 522
 - RSR tracking type 522
 - share level 522
- partitions
 - automatic definition 302
 - bit map block 96
 - changing 404
 - boundaries 406
 - key ranges 406
 - overview 404
 - partition structure modification 405
 - changing with the Partition Definition utility (PDU) 300
 - creating with the Partition Definition utility (PDU) 299
 - data sets, maximum 303
- partitions (*continued*)
 - defining
 - automatically 302
 - manually 302
 - high key 301
 - high key value 301
 - manual definition 302
 - modifying 404
 - boundaries 406
 - key ranges 406
 - overview 404
 - partition structure 405
 - naming conventions 23
 - offline reorganization
 - reallocating data sets 368
 - reloading 369
 - unloading 367
 - updating ILDS 369
 - partition definition process 299
 - partition high key 301
 - reallocating data sets
 - offline reorganization 368
 - reloading
 - offline reorganization 369
 - unloading
 - offline reorganization 367
 - updating ILDS
 - offline reorganization 369
- PASSWD parameter 33
- password protection 33
- paths
 - full duplex 482
 - half duplex 482
 - in hierarchy 8
 - in logical relationships 168
 - third access 482
- PCB (program communication block)
 - alternate PCB statement 306
 - coding 305
 - database PCB statement 306
 - introduction 18
 - maximum number of database PCBs in a PSB 306
- PCF (physical child first) pointers
 - correcting 515
 - description 88
- PCL (physical child last) pointers
 - correcting 515
 - description 89
- PDS directory 520
- performance
 - avoiding split segments 220
 - CI reclaim for KSDSs 241
 - comparison of databases 82
 - discussion 245, 271
 - HISAM 74, 78
 - HSAM 68
 - logical relationships 189
 - monitoring 341
 - tuning a database 347
- PHDAM
 - adjusting 411

- PHDAM (*continued*)
 - options, adjusting 411
- PHDAM (partitioned Hierarchical Direct Access Method)
 - RAPs (root anchor points) 97
- PHDAM (Partitioned Hierarchical Direct Access Method)
 - access methods 11
 - accessing segments 103
 - calls against 84
 - changing DL/I access methods
 - from HDAM 401
 - to HDAM 402
 - counter area of segments, introduction 15
 - data set naming conventions 23
 - database
 - reorganizing 364
 - database records 100
 - database records, locking 109
 - DBCTL support 60
 - deleting segments 107
 - description of 82
 - format of database 95
 - inserting segments 104
 - loading the database 337
 - locking 111
 - logical record length 252
 - multiple data set groups 237
 - options available 84
 - overflow area 98
 - pointer area of segments, introduction 15
 - pointers in 85
 - randomizing module 247
 - root addressable area 98, 100
 - segment format 100
 - size of root addressable area 246
 - space calculations 315
 - specifying free space 245
 - storage of records 98
 - z/OS access methods used 83
- PHIDAM
 - access methods 11
- PHIDAM (Partitioned Hierarchical Indexed Direct Access Method)
 - accessing segments 103
 - calls against 84
 - changing DL/I access methods
 - from HIDAM 401
 - to HIDAM 402
 - counter area of segments, introduction 15
 - data set naming conventions 23
 - database
 - reorganizing 364
 - DBCTL support 60
 - deleting segments 107
 - description of 82
 - format of database 95
 - index database 83, 100
 - index segment 102
 - inserting segments 104
 - loading the database 337
 - locking 111
 - logical record length 252
- PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) (*continued*)
 - maximum size 83
 - multiple data set groups 237
 - options available 84
 - pointer area of segments, introduction 15
 - pointers in 85
 - segment format 101
 - space calculations 315
 - specifying free space 245
 - storage of records 100
 - when to use 85
- PHIDAM (Partitioned Indexed Hierarchical Indexed Direct Access Method)
 - space calculations 109
- physical block size 252
- physical child first pointers 88, 515
- physical child last pointers 89, 515
- physical parent in logical relationships 158, 162
- physical parent pointer
 - See PP (physical parent) pointer 165
- physical twin backward pointers 92, 515
- physical twin forward pointers 91, 515
- physically adjacent 64, 69
- PI (program isolation), lock protocols 109
- pointer area
 - of segment, introduction 15
- pointer field 200
- POINTER parameter 181
- pointer segment 194, 199
- pointers
 - correcting 515
 - direct-address 82
 - FCP (forward chain pointer) 134
 - HALDB self-healing pointer process 388
 - performance 392
 - HB (hierarchic backward) 87
 - HD 85
 - hierarchic forward (HF) 86
 - HISAM (Hierarchical Indexed Sequential Access Method) 71
 - in logical relationships 167
 - in secondary indexes 200, 201
 - LCF 164
 - LCL 164
 - logical relationships 162
 - logical twin 515
 - LP (logical parent) 162, 515
 - LTB 166
 - LTF 166
 - mixing types 93
 - PCF (physical child first) 88
 - PCL (physical child last) 89
 - PP 165
 - PTB 92
 - PTF 91
 - self-healing pointer process 388
 - performance 392
 - sequence in a segment's prefix 94, 170
 - symbolic 195, 200
 - types 397

- position
 - hierarchy 10
 - MSDB 137
- post-implementation review 29
- PP (physical parent) pointer 165
- pre-formatting data set space 267
- preallocated CIs 274
- prefix descriptor byte 469
- prefix part of segment 14
- Prefix Resolution utility (DFSURG10) 357
- Prefix Update utility (DFSURGP0) 358
- preopen
 - disabling for DEDB areas 116
- preopening
 - DEDB areas 115
- Prereorganization utility (DFSURPR0) 356
- primary data set, defined 69
- private buffer pool
 - description 144
- procedures
 - adding logical relationships 433
 - adding secondary indexes 451
 - adding segment edit/compression exit routine 452
 - adding segment types 430
 - adding variable-length segments 451
 - Asynchronous Data Capture 453
 - calculating database size 315
 - changing DASD 409
 - changing hierarchic structure
 - changing sequence of segment types 408
 - combining segments 408
 - changing segment size 432
 - converting concatenated keys 454
 - deleting segment types 431
 - description of 19
 - introduction 6
 - modifying a database 429
 - reorganization
 - HD database 364
 - HISAM database 364
 - primary index 364
- processing option H 285
- processing option P
 - and NBA limit 289
 - and NBA/FPB limit 293
 - in determining the size of the UOW 275
- processing, mixed mode 131
- PROCOPT parameter
 - establishing security 32
 - in HSSP 285
 - option H 285
 - option K 307
 - option P 275
- PROCSEQ parameter 194, 197
- program communication block
 - See PCB (program communication block)
- program communication block (PCB)
 - database PCB statement 306
 - maximum number of database PCBs in a PSB 306
- program isolation lock manager 109

- program specification block
 - See PSB (program specification block)
- programs
 - DB Monitor 341
 - DB Monitor Report print 341
 - DFSDDLTO 314
 - DL/I test 314
 - IEFBR14 utility 322
 - IEHPROGM program 322
 - running 333
 - writing a load program 325, 336
- PROT parameter 206
- PSB (program specification block)
 - as mask over data structure 31
 - coding 305
 - defined 18
 - using dictionary to generate 18
- PSBGEN (Program Specification Block Generation) 307
 - utilities 305, 460
- PSBLIB library 305
- PSINDEX
 - data set naming conventions 23
 - database
 - reorganizing 364
 - DDNAME requirements 23
- PTB (physical twin backward) 515
- PTB (physical twin backward) pointers 92
- PTF (physical twin forward) 515
- PTF (physical twin forward) pointers 91

Q

- Q command codes, locking 110
- QSAM (Queued Sequential Access Method)
 - access to GSAM databases 80
 - and OSAM data set 513
 - processing HSAM databases 65
 - processing SHSAM databases 79

R

- randomizing module
 - DEDB design 275
 - in HDAM database records 247
 - in PHDAM database records 247
 - introduction 83
- randomizing routine
 - changing online 458
 - deleting 458
- randomizing routines
 - DEDB, standard
 - changing RAP space allocation online 464
 - DEDB, two stage
 - changing online 460
 - introduction to changing online 457
- RAP (root anchor point)
 - changing randomizing routines 457
- RAPs (root anchor points)
 - explained 97
 - HIDAM 102

- RAPs (root anchor points) *(continued)*
 - number 99
- RATE parameter of INITIATE OLREORG
 - command 383
- RDF (record definition field) 318
- read errors
 - DEDB
 - VSO 151
- real logical child 161, 164, 192
- RECON data set
 - selecting for HALDB (High Availability Large Database) 545
- record deactivation 118
- Record Deactivation 118
- record definition field (RDF) 318
- record distribution
 - DEDB
 - impact of changes to UOW structure on distribution 463
- RECORD parameter 252
- record search argument (RSA) 80
- Recoverable Resource Manager Services attachment facility 61
- recovery 268
 - overview 5
- recovery for HALDB Online Reorganization 386
- recursive structures 172, 176, 214
- registering databases 155
- relative block number 99
- reload utility (DFSURGL0) 355
- reload utility (DFSURRL0) 354
- Remote Site Recovery (RSR)
 - DBTRACK 522
 - global service group
 - HALDB (High Availability Large Database) 522
 - HALDB Online Reorganization 384
 - RCVTRACK 522
 - shadowing
 - HALDB (High Availability Large Database) 522
 - tracking type
 - HALDB (High Availability Large Database) 522
- reorganization 349
 - online
 - HALDB naming convention 378
- reorganization utilities
 - See also* utilities
 - introduction to reorganization utilities 349
- reorganizing 347, 515
 - assessing need using Database Surveyor utility 361
 - Database Surveyor utility (DFSPRSUR) 361
 - HALDB (High Availability Large Database) 364
 - offline reorganization 365
 - overview of offline reorganization 365
 - reallocating data sets 368
 - reloading partitions 369
 - secondary indexes 370
 - unloading partitions 367
 - updating ILDS 369
 - HALDB self-healing pointer process 388
 - offline reorganization
 - HALDB (High Availability Large Database) 365
 - reorganizing *(continued)*
 - offline reorganization *(continued)*
 - reallocating data sets 368
 - reloading HALDB partitions 369
 - unloading HALDB partitions 367
 - updating ILDS 369
 - PHDAM database
 - overview of offline reorganization 365
 - PHDAM databases 364
 - PHIDAM database
 - overview of offline reorganization 365
 - PHIDAM databases 364
 - reloading HALDB partitions 369
 - secondary indexes
 - HALDB (High Availability Large Database) 370
 - self-healing pointer process for HALDBs 388
 - unloading HALDB partitions 367, 368
 - updating ILDS 369
- REPL parameter 228
- replace rules
 - AM status code 475
 - DA status code 475
 - RX status code 475
 - status codes 475
- replace rules for logical relationships
 - choosing 189
 - description of 475, 479
- replacing segments
 - HISAM databases 78
 - HSAM databases 68
- REPLICATE | NOREPLICATE parameter 268
- reports
 - Fast Path Analysis 345
- resolution utility (DFSURG10) 357
- resolving data conflicts 56
- resource allocation for MSDBs 279
- resource contention 280
- restart 80
 - emergency
 - reopening DEDB areas 115
 - HALDB Online Reorganization 383, 384
- restrictions
 - HALDB Online Reorganization 376
 - HSSP, of 284
 - modifying existing logical relationships 449
 - segments 14
 - using secondary indexes with logical relationships 209
- reviews 25
- RMNAME parameter 248
 - specifying number of blocks or CIs 247
 - specifying number of RAPS 97
 - usage 457
- ROLB call 288, 292
- root addressable area 98, 123
- root addressable space
 - changing online, with two stage randomizing routine 460
- root anchor point (RAP)
 - changing randomizing routines 457

- root anchor points
 - See RAPs (root anchor points) 97
- root processing
 - sequential
 - HIDAM 103
- root segment, definition 7
- RRSAF
 - See Recoverable Resource Manager Services
 - attachment facility
- RSA (record search argument) 80
- rules
 - defining logical relationships 182
 - description of 471, 511
 - in logical databases 183, 189
 - in physical databases 181
 - fields in a segment 15
 - HD with data set groups 237
 - secondary indexes with logical relationships 209
 - segments 14
 - sequence fields 16
 - using an SSA 135
- RULES parameter 471, 511
- RX status code 475

S

- SB (OSAM Sequential Buffering)
 - benefits 258
 - productivity 258
 - programs 258
 - utilities 258
 - buffer handler 260
 - buffer pools 260
 - buffer set 260
 - CICS 258
 - conditional activation 259
 - data set groups 259
 - DB-PCP/DSG pair 259
 - deactivation 259
 - description 257, 258
 - disallowing use 263
 - HALDB Online Reorganization 388
 - overlapped I/O 258, 260
 - periodical evaluation 259
 - random read 257
 - requesting use 261, 264
 - sequential read 257
 - virtual storage 260
- scan utility (DFSURGS0) 356
- SCD (system contents directory) 136
- SDEP (sequential dependent)
 - CI preallocation 274
- SDFSRESL 459
- search field 200
- secondary data structure 198
- secondary indexes
 - HALDB (High Availability Large Database)
 - reorganizing 370
 - reorganizing
 - HALDB (High Availability Large Database) 370

- secondary indexing
 - analyzing requirements 56
 - comparison with logical relationships 214
 - description of 192
 - index maintenance exit routine 204
 - INDICES parameter 207
 - introduction 17
 - loading databases 337
 - locking 111
 - maintenance 205
 - making keys unique 202
 - pointer segment 199
 - procedure for adding 451
 - processing as separate database 206
 - restructured hierarchy 197
 - segments 194
 - sharing 207
 - sparse indexing 204
 - specifying in DBD 211
 - storage 198
 - suppressing index entries 204
 - system related fields 202
 - use
 - logical relationships 209
 - variable-length segments 210
 - uses 192
 - utility unload 359
- secondary processing sequence 198
- security
 - establishing 31
 - field-level sensitivity 226
 - introduction 6, 18
- security inspection 29
- SEGM statement 181
 - description 297
 - example 183
 - in secondary indexing 213
 - in the physical DBD 178
 - specifying insert, delete, and replace rules 471
 - specifying variable-length segments 215
- segment
 - data
 - compressing 218
 - editing 218
- segment code
 - description 14
 - HDAM 100
 - HISAM 70
 - HSAM 66
 - PHDAM 100
- Segment compression routine
 - adding 458
 - changing 458
 - deleting 458
- segment deletion 131
- segment edit/compression exit routine
 - avoiding split segments 220
 - introduction 17
 - procedure for adding 452
 - specifying minimum segment size 220
 - specifying the use of 220

- Segment Edit/Compression exit routine
 - description of 218
 - uses 218
- Segment Edit/Compression facility
 - See See segment edit/compression exit routine
- SEGMENT parameter 211
- segment search argument
 - See SSA (segment search argument) 201
- segments
 - accessing
 - HDAM databases 103
 - HIDAM databases 103
 - HISAM databases 72
 - HSAM databases 67
 - PHDAM databases 103
 - PHIDAM databases 103
 - adding to DEDB 462
 - calculating frequency 316
 - calculating size 315
 - changing position of data 433
 - changing size 432
 - child, definition 7
 - counter area 15
 - data elements 15
 - DEDB
 - segment growth 220
 - definition 6
 - deleting
 - HD databases 107
 - HISAM databases 76
 - HSAM databases 68
 - MSDB (main storage database) 136
 - deleting from DEDB 462
 - dependent, definition 7
 - fields 15
 - fixed-length 14
 - fixed-length segments
 - specifying minimum size 220
 - full-function
 - avoiding split segments 220
 - specifying minimum size 220
 - inserting
 - HD databases 104
 - HISAM databases 72
 - HSAM databases 68
 - MSDB 136
 - introduction to 14
 - logical child 169
 - moving segment types 432
 - occurrence, definition 7
 - parent, definition 7
 - pointer 194
 - pointer area 15
 - procedure for adding to database 430
 - procedure for deleting from database 431
 - replacing
 - HISAM databases 78
 - HSAM databases 68
 - root, definition 7
 - rules 14
 - source 195

- segments (*continued*)
 - target 195
 - twin, definition 8
 - type, definition 7
 - variable length 14
 - variable-length 215
 - variable-length segments
 - specifying minimum size 220
- self-healing pointer process 388
 - performance 392
- SENFLD statement 226, 307
- SENSEG statement
 - description 307
 - field-level sensitivity 227
 - restricting data access 31
- sequence field
 - See also keys
 - HIDAM 101
 - HISAM 69
 - HSAM (Hierarchical Sequential Access Method) 65
 - introduction to 15
 - logical relationships 176, 177
 - PHIDAM (Partitioned Hierarchical Indexed Direct Access Method) 101
 - unique, definition 16
- sequence set records 268
- sequencing in hierarchy 9
- sequencing logical twin chains 191
- sequential access methods
 - HISAM 69
 - HSAM 64
- sequential buffering (SB)
 - See SB (OSAM Sequential Buffering) 257
- sequential dependent part of area 123
- sequential randomizing module 247
- sequential root processing
 - HIDAM 103
- sequential storage method 60
- SETO statement 285
- SETR statement 285
- shared secondary indexes 207
- SHARELVL 120
- SHISAM (Simple Hierarchical Indexed Sequential Access Method) 78, 337
 - CI reclaim restriction 241
 - VSAM REPRO, using 241
- SHSAM (Simple Hierarchical Sequential Access Method) 78, 79
- Simple Hierarchical Indexed Sequential Access Method (SHISAM)
 - See SHISAM (Simple Hierarchical Indexed Sequential Access Method) 78
- Simple Hierarchical Sequential Access Method (SHSAM)
 - See SHSAM (Simple Hierarchical Sequential Access Method) 78
- single area data sets (ADS)
 - Fast Path I/O toleration 154
 - I/O errors 154

- size
 - maximum
 - HALDB (High Availability Large Database) 83
 - HIDAM database 83
 - PHDAM database 83
 - PHIDAM database 83
- size calculations
 - See space calculations 315
- size field in variable-length segments 215
- size of DEDB estimation 274
- SOURCE parameter 181, 190
- source segment 195
- space calculations
 - CIs or blocks needed for database 318
 - database size 315
 - overhead for DEDB CI resources 317
- space management fields, updating 105
- space management in HD databases 95
- space release in logical relationships 484
- space search algorithm
 - HD databases 107
- sparse indexing 204
- SPEED I RECOVERY parameter 267
- SRCH parameter 212
- SSA (segment search argument)
 - restrictions for DEDBs 131
 - secondary indexes 201
- standards and procedures
 - description of 19
 - introduction 6
- START parameter 203
- statements
 - AREA 296
 - data set
 - description of 296
 - DATASET
 - example of 240
 - specifying DDNAMEs for data sets 183
 - DBD 213, 296
 - DBDGEN 298
 - END 298, 307
 - FIELD
 - definition of 202
 - in the DBD 269
 - position in DBD 297
 - FINISH 298
 - LCHILD in logical relationships 178, 181, 211, 297
 - OPTIONS
 - fixing buffers in VSAM 256
 - for OSAM 269
 - for VSAM 264, 266
 - OSAM 269
 - use in splitting CIs 73
 - PSBGEN 307
 - SEGM
 - description of 297
 - example of 183, 213
 - in secondary indexing 213
 - in the physical DBD 178, 181
 - specifying insert, delete, and replace rules 471
 - specifying variable-length segments 215
- statements (*continued*)
 - SENFLD 226, 307
 - SENSEG
 - description of 307
 - field-level sensitivity 227
 - restricting data access 31
 - XDFLD
 - description of 202
 - in secondary indexing 211
 - restrictions in use 298
 - specifying sparse indexing 204
- status codes
 - AM
 - in a delete call 483
 - in an insert call 473
 - DA 483
 - DX 483
 - FH 117
 - for replace rules 475
 - FR
 - for BMP regions 289
 - for CCTL threads 293
 - in Fast Path buffer allocation 288
 - in Fast Path buffer allocation for BMPs 292
 - FW
 - for CCTL threads 293
 - in BMP regions 289
 - in Fast Path buffer allocation 288
 - in Fast Path buffer allocation for BMPs 292
 - GC 274
 - GE 177, 473
 - II 473
 - IX 473
 - NE 206
- storage of data
 - DEDBs 126
 - HDAM databases 98
 - HIDAM databases 100
 - HISAM databases 69
 - HSAM databases 65
 - introduction 6
 - MSDB (main storage database) 134, 283
 - multiple data set groups 238
 - PHDAM databases 98
 - PHIDAM databases 100
 - variable-length segments 216
- SUBSEQ parameter 202
- subsequence field 200
- subset pointers 127, 277
- suppressing index entries 204
- Surveyor utility (DFSPRSUR) 361
- SX (/SX) operand 202
- symbolic checkpoint call 80
- symbolic pointers
 - logical relationships 163, 190
 - secondary indexes 195, 201
- SYNC (Synchronization Point) call 274
- sync point processing for Fast Path 153
- synchronization point
 - Fast Path 153, 289, 293
 - output thread 153

- synchronization point (*continued*)
 - processing 153, 425
- synonyms 100
- syntax diagram
 - how to read xxvi
- system contents directory (SCD) 136
- system related fields 202

T

- tape, magnetic 64
- target segment 195
- task ID field 97
- terminal-related database 132
- termination phase of HALDB Online
 - Reorganization 374
- test databases 311
 - developing
 - Cross System Product 313
 - Data Extraction, Processing, and Restructuring System 313
 - DL/I test program (DFSDDLTO) 314
 - File Manager for z/OS for IMS Data 314
 - IMS Application Development Facility II 313
 - VisualAge Generator 313
- testing a database
 - introduction 5
- testing, application programs 312
- third access path 482
- tools
 - Data Extraction, Processing, and Restructuring System 313
 - for test databases 313
 - DL/I test program 314
 - VisualAge Generator 313
- trace parameters
 - using 266
- track space used 252
- transaction timings, Fast Path 344
- tuning a database
 - description of 347
 - Fast Path 343
 - introduction 5
- two stage randomizing routine
 - changing root addressable space 460
- TYPE parameter 228

U

- UCF (utility control facility)
 - described 361
 - restartable initial database load program 332
 - running restartable load program under 333
- unique sequence fields
 - HISAM (Hierarchical Indexed Sequential Access Method) 69
 - introduction 16
- unit of reorganization for HALDB Online
 - Reorganization 373
- units of work (UOW) 123
- Unload utility (DFSURGU0) 354

- unload utility (DFSURUL0) 353
- UOW (unit of work) 123, 274
- UOW locking 286
- UOW structural definition 460
- use chain 253
- user data field in pointer segment 202
- utilities
 - Application Control Block Maintenance utility
 - building the ACBs 308
 - Database Change Accumulation 387
 - database image copy 387
 - Database Prefix Resolution utility (DFSURG10) 357
 - Database Prefix Update utility (DFSURGP0) 358
 - Database Preorganization utility (DFSURPR0) 356
 - Database Scan utility (DFSURGS0) 356
 - Database Surveyor (DFSPRSUR) 361
 - DBDGEN 295
 - DBFDBMA0 133
 - DBFUHDR0 274
 - DFSPRCT1 362
 - DFSPRSUR 361
 - DFSUCF00 361
 - DFSURG10 357
 - DFSURGL0 355
 - DFSURGP0 358
 - DFSURGS0 356
 - DFSURGU0 354
 - DFSURPR0 356
 - DFSURRL0 354
 - DFSURUL0 353
 - for unload and reloading secondary indexes 359
 - HALDB Online Reorganization 385
 - HD Reorganization Reload 355
 - HD Reorganization Unload 354
 - High-Speed DEDB Direct Reorganization (DBFUHDR0) 274
 - HISAM Reorganization Reload 354
 - HISAM Reorganization Unload 353
 - MSDB Maintenance 133
 - Partial Database Reorganization 362
 - PSBGEN 305
 - reorganization 349
 - UCF 361
 - Unload 354
- utility control facility
 - See UCF (utility control facility)

V

- variable intersection data (VID) 171
- variable-length segments
 - definition 14
 - description of 215
 - introduction 17
 - procedure for adding 451
 - replace operations 216
 - specifying in DBD 215
 - specifying minimum size 220
 - storage 216
 - use with secondary indexes 210

- variable-length segments *(continued)*
 - uses 217
 - using 215
 - what application programmers need to know 217
- VERSION parameter 222
- VID (variable intersection data) 171
- virtual logical child 161
- Virtual Storage Access Method (VSAM)
 - HISAM databases 69
- virtual storage option
 - introduction 139
- VisualAge Generator 313
- VSAM
 - data sets
 - maximum size 83
- VSAM (Virtual Storage Access Method)
 - access to GSAM databases 80
 - adjusting buffers 411
 - adjusting options 415, 416
 - and Hiperspace buffering 254
 - changing access methods 417
 - changing space allocation 416
 - CIDF (control interval definition field) 318
 - ESDS in HD databases 95
 - HISAM databases 69
 - index 268
 - local shared resource pools
 - assigning data sets 266
 - defining 266
 - index and data subpools 266
 - subpools of same size 254
 - options 264, 269
 - passwords 33
 - RDF (record definition field) 318
 - storage of secondary indexes 198
 - track space used 252
- VSAMFIX parameter 256, 266
- VSAMPLS parameter 266
- VSO DEDB (virtual storage option data entry database)
 - checkpoint processing 152
 - data sharing 149
 - defining a VSO DEDB area 140
 - defining a VSO DEDB cache structure name 144
 - emergency restart 152
 - I/O error processing 151
 - read errors 151
 - write errors 151
 - input processing 150
 - locking 148
 - options across restart 152
 - output processing 150
 - PRELOAD option 151
 - resource control 148
 - using data spaces 147
 - with XRF 153
- VSO DEDB areas
 - block-level sharing of 142
 - defining
 - CHANGE.DBDS 139
 - INIT.DBDS 139

- VSO DEDB areas *(continued)*
 - virtual storage
 - coupling facility cache structure 139
 - data space 139

W

- write errors, DEDB VSO 151

X

- XDFLD statement
 - description 202
 - in secondary indexing 211
 - restrictions in use 298
 - specifying sparse indexing 204
- XML
 - decomposed storage
 - overview 244
 - intact storage
 - overview 244
 - overview of storing in IMS databases 244
 - schema
 - overview of storing XML data 244

Z

- z/OS access methods
 - used by HD 83
 - used by HSAM 65



Program Number: 5655-J38

Printed in USA

SC18-7806-01



Spine information:



IMS

Administration Guide: Database Manager

Version 9