

**Communicator 3000
MPE/iX Release 6.5
(Non-Platform Software Release C.65.00)**

HP 3000 MPE/iX Computer Systems

Volume 11, Issue 1



**Manufacturing Part Number: 30216-90291
E0300**

U.S.A. March 2000

Notice

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability or fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for direct, indirect, special, incidental or consequential damages in connection with the furnishing or use of this material.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

This document contains proprietary information which is protected by copyright. All rights reserved. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

Restricted Rights Legend

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013. Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19 (c) (1,2).

Acknowledgments

UNIX is a registered trademark of The Open Group.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304 U.S.A.

© Copyright 1997 - 2000 by Hewlett-Packard Company

Contents

1. Announcements

Introducing MPE/iX 6.5 Release	11
Removal of Support for Servers and HP-IB in MPE/iX	12
Discontinuance of High Availability Fiber Link Disk Drives	13
Obsolete Products Removed from Release 6.5	14
During an FOS Update.	15
Following an Install	15
Files Associated with Obsolete Products:	
Resource Sharing, CCSY Access Server, CCSY Access SQL, Cooperative Services/XL	16
Important Support Changes in MPE/iX 6.5	19
Product Overview	19
Predictive Support Changes for MPE/iX Release 6.5	19
Obtaining Software Security Patches for your HP Computer System	22
HP Electronic Support Center Security-Related Services	22
Subscribing to HP Electronic Support Center Security Bulletin Services	23
Instructions	23
If You Discover a Security Problem	23

2. Overview—MPE/iX Release 6.5 (C.65.00)

Communicator Summary	26
Chapter 1 Announcements	26
Chapter 2, Overview--Communicator Summary.	26
Chapter 3, Growth Solution	26
Chapter 4, DCM	27
Chapter 5, Internet, Interoperability and Databases	27
Chapter 6, Announcing a Post-6.5 Release Patch	28
Chapter 7, How to Order	28
Chapter 8, Product Release History.	28
Chapter 9, Catalog of User Documentation.	28
MPE/iX Patches on HP IT Resource Center	29
Features and Benefits	29
Access Method to the HP Electronic Support Center	29
Electronic Digests	30
Patch Installation Tools	30

3. Growth Solution

MPE/iX Large Files Overview	34
Introduction.	34
Summary of Large File Compatibility Issues	35
FLABELINFO	37
32-bit Integer Overflow.	37
AIF:OS Impact	39

Contents

OPTIONS Enhancement	41
Command Interpreter Changes	41
Intrinsics Changes	46
New Intrinsics	48
Miscellaneous	52
AIFs for Large Files.	55
Enhancement for AIFFILEGGET/AIFFILEGPUT	55
Enhancement for AIFFILELGET/AIFFILELPUT	55
Error Messages	56
AIFs for Multiple Job Queues	57
Enhancement for AIFSYSWIDEGET	57
Enhancement for AIFJSGET	57
Increase in Number of TCP Connections	58
Introduction	58
Configuration through NMMGR	58
Dispatcher Multiple Runqueues	60
Introduction	60
Technical Overview	60
Benefits to Customers	60
Conclusion.	61
Increased System Limits - An Overview	62
Support for 511 Disks on MPE/iX.	63
Introduction	63
Configuring More Than 255 Disks	63
User Interface Changes	63
Forward/Backward Compatibility Issues	63
Increased Number of DSTs	64
Introduction	64
Technical Overview	64
Benefits to Customers	64
Conclusion.	64
User-Defined Job Queues and ALTJOB HIPRI Enhancements	66
Enhancements Summary.	66
New Commands	67
Modified Commands.	69
NEWCI Command (CI Elimination).	74
Background Information	74
Operation and Use of New Command.	76
Compatibility Considerations for COBOL and C.	79
Forward Compatibility Not a Problem	79
Backward Compatibility Not Guaranteed	79
COBOL/iX.	80

Contents

C/iX	80
C/iX Supports 64-bit Integers	82
FTP/iX Support for Large Files	84
Features	84
User Interface	84
Installation	85

4. Data Center Management

Enterprise Management Solution	88
The Current Need	88
Recent Offering	88
Short Term Plans	89
Long Term Plans	89
Summary	90
Support Tools Manager (STM)	91
Enhanced Message Source Template for HP 3000 Agents in ITO	93
Introduction	93
OS Monitoring	93
Disk Drive Monitor	94
Volume Set Monitor	94
Job and Process Monitor	94
Monitor for other Monitors	94
Printer Outfence Monitor	94
Printer Queue Length Monitor	95
Pending Console Reply Monitor	95
Message Queue File Monitor	95
Networking Monitoring	95
Networking Interface Monitoring	95
Network Services Monitoring	96
Database Monitoring	96
Console Messages	96
Volume Management	96
Security	96
Backup/Tape Volumes	97
Open View System Manager	97
HP Secure Web Console	98
Product Overview	98
Installation	98
Product Structure	99
Product Features	99
Support from the World Wide Web	100
TurboSTORE Support for DLT	101

Contents

About the DLT Drives	101
TurboSTORE Products and the DLT.	101
About Fast Search on Tapes.	101
Fast Search on DLT Tape Drives.	101
Limitation of JXX0.	102
DLT4000/DLT7000 Differential Tape Support on MPE/IX.	103
Software Requirements	103
Hardware Requirements	103
Product Overview.	104
Configuration	106
Announcing DLT8000 Differential Tape Support on MPE/iX.	107
Software Requirements	107
Hardware Requirements	107
Product Overview.	108
Configuration	108
TERMDSM: User Interface Change under System Diagnostics.	110
What It Means to the User	110
What the User Will See	110
Other Commands of Interest	112
HP SureStore E Disk Array 12H on the HP 3000	113
Site, Software and Hardware Requirements	113
AutoRAID Configuration Requirements.	113
Host Hardware/Firmware Requirements	113
Host Software Requirements.	114
Host Diagnostic Support	114
Off-line Diagnostics	115
Configuring the 12H for Performance on the HP 3000	115
DVD Supported on the HP 3000.	116
Product Numbers	116
DLT7000 and Legato Storage Node Usage.	117
HP SureStore E Disk Array XP256	118
Announcing Support of the HP SureStore E Disk Array XP256.	118
SureStore E Disk Array XP256 Business Copy XPand/or Continuous Access XP in an HP3000 Environment 120	
Product Descriptions	120
Environment.	120
Basic Operation	120
System A Procedures	121
System B Procedures	122
Business Copy Procedures	123

5. Internet and Interoperability

Contents

Large Transactions for IMAGE Users	126
Background	126
Current Limitations	126
XM User Checkpoint Priority Control	131
Background	131
Solution Objective	131
Overview of the Solution	131
Transact/iX Expanded B-Tree and File Open Enhancement	132
Support of B-Trees in Transact	132
File Open Transact	135
Apache for MPE/iX 6.5	138
Overview	138
Feature Set	138
Technical Overview	138
Major Components	139
Introducing LDAP C-SDK/iX	145
MPE/iX Software Developer's Kit (SDK) for Java, Version 1.2.2	146
HP Driver for JDBC	148
HP JDBC Components	149
Requirements	149
Installation	150
Using the HP Driver for JDBC	153
ALLBASE/SQL Specifics	155
Troubleshooting	159
6. Announcing a Post 6.5 Release Patch	
IMAGE/SQL Date Mapping Enhancement	164
SP(LIT)	165
U[PCDATE] TYPE	166
Format(TI)Type StorageType #Bytes Explanation	167
Notes	167
Examples	168
7. Product List—How to Order	
8. Product Release History	
Product Changes by Releases	173
Supported Releases	176
9. Catalog of User Documentation	
MPE/iX 6.5 New or Updated Manuals	182
Manual Collections	184

Tables

Table 5-1.. Data Type Mapping	156
Table 5-2.. Data Type Conversions	157
Table 5-3.. Unsupported Statements	158
Table 8-1.. MPE/iX Product Releases	173
Table 8-2.. Supported System Release Matrix	177
Table 9-1.. MPE/iX 6.5 New or Updated Manuals	182
Table 9-2.. Manual Collections	184

1 Announcements

Introducing MPE/iX 6.5 Release

Welcome to the exciting MPE/iX 6.5 release. New releases of MPE/iX usually have a short set of focus areas where we have concentrated our efforts to have a major impact. In the case of MPE/iX 6.5, those areas of focus are performance scaling and capacity.

The main driver for the emphasis on performance scaling and capacity improvements is that HP 3000 customers were growing at a rate that is starting to tax the performance of our most powerful servers. For these customers to take full advantage of new hardware which we will be supporting in the HP 3000 family, we needed to address certain performance bottlenecks.

In the MPE/iX 6.5 release, we have made major changes to the core operating system, while improving the performance of HP 3000 networking and IMAGE/SQL performance. Some of these changes will only be apparent on those high-end systems used by customers already pushing the performance of those systems. However, these changes will benefit many more customers by ensuring a future upgrade path and helping to build a bright future for the HP 3000.

Of course, with the added performance, comes the reality of more users and larger workloads. All of this stresses the limits and capacity of the system in many ways. So we have also greatly enhanced the capacity and limits of the HP 3000, adding support for:

- Large Files (>4GB)
- Increased disk spindle limits (up to 511 spindles)
- Larger physical memories (up to 16 GB on some systems)
- Larger network limits (TIO connections, network sockets, VT)
- Improved handling of large Image/SQL transactions

In addition, we've also made investments in other key areas. In today's world, interoperability is key to enabling our customers to compete in the fast changing world of e-commerce. So we've continued our investments in key interoperability solutions by fully supporting the very popular Apache web server and LDAP (Lightweight Directory Access Protocol) software on the HP 3000. We've also improved Java Database Connectivity (JDBC) support on the HP 3000. And don't forget, many products and tools continue to be delivered on the HP 3000 separate from major releases, including support for many new peripherals.

Removal of Support for Servers and HP-IB in MPE/iX

In order to streamline the development and test of future MPE/iX releases, the PBA-IB HP-IB Device Adapter (A1747A) and old CIB I/O based HP 3000 Servers will not work with MPE/iX Release 6.5 and later releases.

The PBA-IB HP-IB Device Adapter was discontinued in May 1999 and has an end of support date of November 1, 2001. The following old CIB I/O (also known as CIO) based HP 3000 Servers will have completed their end-of-support well before November 1, 2001:

- the 925/935/949 family
- the 920/922/932/948/958
- the 950/955/960/980 family

HP-IB and the above HP 3000 Servers will continue to be supported through their end-of-support date on MPE/iX Release 5.5 (which has a projected end-of-support date of November 1, 2000) and Release 6.0 (which has a projected end-of-support of November 1, 2001).

NOTE If you are running MPE/iX on one of these older systems, or on a system with any of these older cards or peripherals, you must **not** update that system to Release 6.5.

Customers wishing to replace their HP-IB I/O cards and peripherals should consider SCSI or LAN connected peripherals. Customers wishing to replace the above older HP 3000 Servers as they complete their support life should consider the 9x8 or 9x9KS HP 3000 Servers.

Discontinuance of High Availability Fiber Link Disk Drives

In May 1993, Hewlett-Packard Company introduced a family of High Availability Fiber Link disk arrays. Those Fiber Link (HP-FL) disk arrays were discontinued in August of 1997, along with other standalone Fiber link disk drives. The Fiber Link disk arrays were HP's first RAID devices built for high availability, high performance, high capacity and distances up to 500 meters. Since 1993, HP has embraced new disk array technologies, EMC disk arrays, High Availability Model 10 and 20 disk arrays and AutoRAID.

CSY is also moving to new processor technologies, like the N-class computers. These new computers call for new I/O cards, devices drives and I/O backplanes. MPE/iX is also changing. MPE/iX 6.5 is the first HP 3000 operating system to support these new I/O requirements. Therefore, MPE/iX 6.5 will not support Fiber Link disk I/O system drivers, I/O cards, and Fiber Link disks. The last release of MPE/iX to support HP-FL drivers and disk is MPE/iX 6.0.

The following is a partial of products not carried forward in MPE/iX 6.5:

Part #	Description of Obsolete Part
C2258HA	1/02 High Availability FL disk array
C2254HA	4/99 High Availability FL disk array
C2252B	4/99 High Availability FL disk array
C2258B	1/02 High Availability FL disk array
C2252HA	4/00 High Availability FL disk array
C2254B	4/99 High Availability FL disk array
C2259B	1/02 High Availability FL disk array
C2259HA	1/02 High Availability FL disk array
C2201A	8/97 FL disk
C2204A	8/97 FL disk
A1748A	PBA FL Chan-span card with HP-FL adapter
A28616A	PBA FL NIO Optic interface card
A27115A	CIO Fiber Optic interface card

Obsolete Products Removed from Release 6.5

by Allan Hertling
CS R&D

Numerous products that have reached their “End of Support Life” were removed from the SUBSYS tape for the 6.5 release.

The following obsolete products have been removed.

Product Number	Product Description
B1710A	NWOffice (FULL)
B1711A	NWOffice (DESK)
B1712A	NWOffice (ACCESS)
B1713A	NWOffice (CORE)
B1714A	NWOffice (FULL)
B1715A	NWOffice (DESK)
B1716A	NWOffice (ACCESS)
B1716B	Access Server/XL
B1717A	NWO-AccessSQL
B1717B	NWO-AccessSQL
B1718A	Resource Sharing
B3160A	CCSY Access Server
B3162A	CCSY Access SQL
32560A	PSP (FULL)
32561A	PSP (DESK)
32562A	PSP (ACCESS)
32563A	PSP (CORE)
32571A	Cooperative Services/XL
32586A	Information Access
35460A	PC Backup/XL
36561A	HPSpell (AM)
36562A	HPSpell (DUAL)
36576A	HPSlate
36930A	DSG-V
50700A	LaserRX

This product removal from the SUBSYS tape will affect you if one or more of these products reside on your system.

During an FOS Update

If your system contained any of the above listed products, the associated library modules were removed from the system libraries. The product files will not be removed.

NOTE The products listed above are obsolete and no longer supported. There is no guarantee that they will continue to function correctly under Release 6.5 or future versions of the operating system.

Following an Install

None of the products listed above will reside on your system after you have completed an `INSTALL` using the Release 6.5 tapes. If any of the products listed above are required on the system, the product files will have to be retrieved from a previous system backup which contains the product files.

NOTE Once again, the products listed above are obsolete and no longer supported. There is no guarantee that they will continue to function correctly under Release 6.5 or future versions of the operating system.

Files Associated with Obsolete Products: Resource Sharing, CCSY Access Server, CCSY Access SQL, Cooperative Services/XL

by Sue Meloy
CS-R&D

The following products are obsolete and no longer being shipped:

- B1718A Resource Sharing
- B3160A CCSY Access Server
- B3162A CCSY Access SQL
- 32571A Coop. Services/XL

Files associated with these products may still exist on your system. The following files are installed by these products and can be removed if you wish to clean up your system. Since some of these files may be accessed by networking software, stop network services on the system before removing the files.

After removing the files, stream JCONFJOB.NET.SYS to rebuild the network services configuration files, then restart the networking software.

Files related to B1718A Resource Sharing	
CDM000.PPC.SYS	JROOTCON.PPC.SYS
CPDMG000.PPC.SYS	PDSMAST.NET.SYS
CRCDG000.PPC.SYS	PDSSERV.NET.SYS
DISCMGR.PPC.SYS	PDSFILE.NET.SYS
RESMGR.PPC.SYS	RSUDC.PPC.SYS
PDBAKUP.PPC.SYS	RSCONF.PPC.SYS
DISCSKAN.PPC.SYS	README.PPC.SYS
PDSERVER.PPC.SYS	NOTICE.PPC.SYS
PDDIAG.PPC.SYS	CLC000.PPC.SYS
ROOTCONV.PPC.SYS	NWOLC.PPC.SYS
NBMON.PPC.SYS	CPDDG000.PPC.SYS
PDXL.PPC.SYS	QDIAG000.PPC.SYS
TRACEMAP.PPC.SYS	PPCDIAG.PPC.SYS
JPDINSTL.PPC.SYS	PPCDIAGC.PPC.SYS
JDISCAN.PPC.SYS	

Also, modules HP32589_01 and NWOLC_01 in XL.PUB.SYS and segment PPCDIAG in SL.PUB.SYS are installed by this product.

Files Related to B3160A CCSY Access Server	
HDPMC000.PPC.SYS	ADMENUS.PPCUTIL.HPOFFICE
HDPHC000.PPC.SYS	DHDPGINP.PPCUTIL.HPOFFICE
ADFRM000.PPC.SYS	DHDPMODS.PPCUTIL.HPOFFICE
TRFRM000.PPC.SYS	DHDPMAP.PPCUTIL.HPOFFICE
HDSP.PPC.SYS	SSAMPL1.PPCUTIL.HPOFFICE
DICINIT.PPC.SYS	SSAMPL2.PPCUTIL.HPOFFICE
DICTCONV.PPC.SYS	SHDPDIC.PPCUTIL.HPOFFICE
HDSPNS.PPC.SYS	SHDPENV.PPCUTIL.HPOFFICE
ADMIN.PPC.SYS	SDIAGDB.PPCUTIL.HPOFFICE
TRANS.PPC.SYS	JDICTCHG.PPCUTIL.HPOFFICE
HDPBATCH.PPC.SYS	JDICTCLN.PPCUTIL.HPOFFICE
HDPUTIL.PPC.SYS	HDSPMAST.NET.SYS
HDPLOAD.PPC.SYS	HDSPFILE.NET.SYS
HDPUNLD.PPC.SYS	HDSPSERV.NET.SYS
ENVINIT.PPC.SYS	ALLDBIN.PPCUTIL.HPOFFICE
XL.PPC.SYS	SAMPL1IN.PPCUTIL.HPOFFICE
WHATSNEW.PPC.SYS	SAMPCMD.PPCUTIL.HPOFFICE
ACUDC.PPCUTIL.HPOFFICE	JINSTHDP.PPCUTIL.HPOFFICE

Files Related to B3162A CCSY Access SQL	
IASOLID.PPC.SYS	IAGRANT.PPC.SYS
IASOLIDV.PPC.SYS	IACAT000.PPC.SYS
IASOLIDC.PPC.SYS	IASQLDAD.PPC.SYS
JIASOLID.PPC.SYS	IASQLSVR.PPC.SYS
IASQLSN2.PPC.SYS	SQLFILE.NET.SYS
IASQL.PPC.SYS	SQLMAST.NET.SYS
IASQLSN1.PPC.SYS	SQLSERV.NET.SYS

Files Related to 32571A Coop. Services/XL	
HCS.PPC.SYS	HCSERVER.PPC.SYS
SHCSAMP.PPCUTIL.HPOFFICE	HCSMAST.NET.SYS
SAMPXEQ.PPCUTIL.HPOFFICE	HCSERV.NET.SYS
JHCSAMP.PPCUTIL.HPOFFICE	HCSFILE.NET.SYS

Important Support Changes in MPE/iX 6.5

by Gary Robillard
CS-R&D

Product Overview

HP Predictive Support provides proactive hardware support and helps increase the uptime of your systems by monitoring system memory and disk/tape drives.

When the HP Predictive Support software detects a potential problem, it sends a message to the HP Response Center. The Response Center portion of the system screens the data and forwards problems requiring further analysis to a Response Center Engineer. If action is needed at your site, the Response Center Engineer and the account Customer Engineer will work with you until the problem is resolved.

This proactive hardware support is provided as part of the HP Hardware and Software Support Agreement.

Predictive Support Changes for MPE/iX Release 6.5

The Predictive Support Monitor (PSMON.PRED.SYS) startup model has been changed. On previous versions of MPE, PSMON was automatically started by the Diagnostic system monitor (DIAGMON.DIAG.SYS), which was part of the sysdiag diagnostic system. The Support Tools Manager (STM) has replaced the sysdiag diagnostic system on MPE iX 6.5. A new job named JPSMON.PRED.SYS was created to start PSMON.

Predictive Support relies on a sysdiag diagnostic library, and because of this, there is currently no way to have PSMON started by the STM diagnostics. PSMON is now launched by streaming the job JPSMON.PRED.SYS. If you want Predictive Support to monitor your system, JPSMON.PRED.SYS must always be running.

WARNING **If JPSMON.PRED.SYS is not running, Predictive Support will not perform its daily scheduled run to monitor your system.**

To determine if JPSMON is running, you can use the MPE iX SHOWJOB command, as follows:

```
:SHOWJOB JOB=PSMON,MANAGER.SYS
```

You should see one active job.

The installation job for Predictive Support will attempt to place a "STREAM JPSMON.PRED.SYS" command in the system startup file (SYSSTART.PUB.SYS), after the STREAMS command. If this fails, then you should manually add a line to SYSSTART.PUB.SYS to stream the JPSMON.PRED.SYS job.

NOTE The sysstart must have a creator of MANAGER.SYS in order to be automatically executed during system startup. Additionally, the first line of sysstart should be the STARTUP directive.

Important Support Changes in MPE/iX 6.5

You should insure that there is a “SPOOLER LP;OPEN” and a “STREAMS 10” command before the “STREAM JPSMON.PRED.SYS” command in the sysstart file. Also recognize that for your system, the line printer might have a different class name than “LP”, and your streams device might have a logical device number other than 10.

Here is an example of what the lines in `sysstart.pub.sys` might look like:

```
Startup
Limit 20,60
Jobfence 7
Spooler lp;start
Streams 10
Continue
Stream jstrtlan.sysmaint.sys
Continue
Stream jpsmon.pred.sys
**
```

HP Predictive supports the following new peripherals:

DGHS04Y	4GB SCSI Disk Drive
ST32272WD	2GB SCSI Disk drive
ST34572N,W,WC,WD	4GB SCSI Disk Drives
ST34573N,W,WC,WD	4GB SCSI Disk Drives
DGHS09Y	9GB SCSI Disk Drive
ST39102LC	9GB SCSI Disk Drive
ST39173N,W,WC,WD	9GB SCSI Disk Drives
ST39175WC,LC,LW	9GB SCSI Disk Drive
ST118273N,W,WC,WD	18GB SCSI Disk Drives
ST318203LC	18GB SCSI Disk Drive
ST318275LC,LW,WC	18GB SCSI Disk Drives
ST136475LW,WC	36GB Disk Drives
DLT4000	SCSI Digital Linear Tape
DLT7000	SCSI Digital Linear Tape

The following peripherals are no longer supported:

Any HPIB or FLEX Interface Disk, including but not limited to:	
HPC2200A	335MB HPIB Disk
HPC2202A	670MB HPIB Disk
HPC2203A	670MB HPIB Disk
HPC2281A	335MB HPIB Disk
HPC2282A	670MB HPIB Disk
HPC2201A	670MB Flex Disk
HPC2204A	1.34GB Flex Disk

The following SCSI disks are no longer supported:

HPC2460x	420MB Disk
HPC2461x	673 MB Disk
HPC2462x	1.36GB Disk
HPC2470x	234MB disk mechanism
HPC2471x	328MB disk mechanism
HPC2472x	422MB disk UPGRADE KIT
HPC2473x	673MB Disk UPGRADE KIT
HPC2474R	1.36GM C3 UPGRADE KIT
HPC3010M1	2GB SCSI Disk (As of April 30, 2000)

Obtaining Software Security Patches for your HP Computer System

Hewlett-Packard would like to make you aware of a special free service provided for all customers of HP 3000 and HP 9000 computer systems. This service gives customers a direct route to Hewlett-Packard for obtaining information relating to the security of their Hewlett-Packard Computer System(s).

Hewlett-Packard issues information on the availability of Software security patches via Security Bulletins to subscribers of the HP Security Bulletin Digest e-mail service, a part of the HP Electronic Support Center. A Hewlett-Packard support contract is NOT required to subscribe to this service to obtain information or security patches. Any purchaser of an HP 3000 or HP 9000 Computer System can make use of the HP Security Bulletin services at no charge.

Customers may also obtain information and Security Bulletin services via the World Wide Web.

A security problem is a software defect that allows unauthorized personnel to gain access to a Computer System or to circumvent any of the mechanisms that protect the confidentiality, integrity or availability of the information stored on the system. When such problems in Hewlett-Packard software are brought to the attention of the company, their resolution is given a very high priority. This resolution is usually in the form of a Security Bulletin which may explain how to correct the problem or describe how to obtain a software security patch that will correct the problem.

Hewlett-Packard has introduced this service as the primary mechanism to alert subscribers to security problems and provide corrections. Hewlett-Packard will not analyze the relevance of any security patch to any individual customer site within the scope of the HP Security Bulletin service. The responsibility for obtaining and applying security patches resides with the customer.

The remainder of this letter outlines the various security related services offered by HP Electronic Support Center and the methods for subscribing to and retrieving information from it. It also outlines how you can inform Hewlett-Packard of potential security concerns you may have with your Hewlett-Packard Computer System.

HP Electronic Support Center Security-Related Services

HP Electronic Support Center offers subscribers the following benefits:

- Receive Security Bulletins automatically when they are published.
- Retrieve the archive list of bulletins issued prior to subscription.
- Download security patches if the subscriber configuration supports it.

Remember, an HP support contract is not required to subscribe to HP Security Bulletin services.

Subscribing to HP Electronic Support Center Security Bulletin Services

Once you have placed your name on the subscriber list for future Security Bulletins (see instructions below), you will receive them via e-mail on the day they are issued by HP.

As referenced below, you can also view a list of past Security Bulletins issued in the “HP Security Bulletins Archive.”

Instructions

To subscribe to automatically receive future NEW HP Security Bulletins from the HP Electronic Support Center via electronic mail, do the following (instructions subject to change without notice):

1. Use your browser to access the HP Electronic Support Center web page at:

http://us-support.external.hp.com US, Canada, Asia-Pacific,
and Latin-America

http://europe-support.external.hp.com Europe

2. Logon with your User ID and password (or register for one). Remember to save the User ID assigned to you, and your password.
3. Once you are on the HP Electronic Support Center home page, click on “Support Information Digests.” On this page, you can subscribe to many different digest services, including the Security Bulletin Digests.

To review Security Bulletins that have already been released, click on “Search Technical Knowledge Base (Security Bulletins only)” on the HP Electronic Support Center home page. Near the bottom of the next page, click on “Browse the HP Security Bulletins Archive.”

Once in the archive, click on “HP-UX Security Patch Matrix” to get a patch matrix of current HP-UX and BLS security patches. Updated daily, this matrix categorizes security patches by platform/OS release, and by Security Bulletin topic.

If You Discover a Security Problem

To report new security vulnerabilities, send e-mail to

security-alert@hp.com

Please encrypt any exploit information using the security-alert PGP key, available from your local key server, or by sending a message with a -subject- (not body) of ‘get key’ (no quotes) to security-alert@hp.com.

Announcements

[Obtaining Software Security Patches for your HP Computer System](#)

2 Overview—MPE/iX Release 6.5 (C.65.00)

This *Communicator 3000* provides general and detailed information on the new and enhanced functionality for the MPE/iX 6.5 Release (C.65.00), as well as information on release strategy and installation prerequisites.

MPE/iX 6.5 has accumulated all of the functionality previously released on MPE/iX 6.0 Express releases such as Year 2000 enhancements, ODBCLink/SE—a 32-bit driver from M.B. Foster Associates, and B-Tree functionality with IMAGE/SQL. In addition, it also offers support for user-defined job queues—for improved system management, SAMBA—for NT integration, Java, FTP enhancements, and many more enhancements.

MPE/iX 6.5 provides the following additional enhancements:

- Support for Large Files (74GB)
- HP Driver for JDBC
- AIFs for Large Files
- AIFs for multiple job queues
- Increased TCP connections
- Support for 511 disks
- Support Tools Manager
- Enterprise Management Solution
- HP Secure Web Console
- Apache for MPE/iX
- LDAP C-SDK/iX
- NEWCI Command
- DLT8000 Differential Tape support

Communicator Summary

Following are brief descriptions of the articles and chapters.

Chapter 1 Announcements

Important announcements regarding availability of products and services are included in this chapter.

- **Introducing MPE/iX 6.5 Release** -- provides overview of 6.5 enhancements.
- **Removal of Support for Servers and HP-IB in MPE/iX** -- details the discontinuance of support for certain older servers and the HP-IB device adapter.
- **Discontinuance of High Availability Fiber Link Disk Drives** -- describes discontinuance of High Availability Fiber Link disk drives.
- **Obsoleted Products Removed from Release 6.5** -- lists the obsolete products that have been removed from the SUBSYS tape for the 6.5 release.
- **Files associated with Obsoleted Products** -- Lists the files associated with these obsoleted products: Resource Sharing, CCSY Access Server, CCSY Access/SQL and Cooperative Services/XL.
- **Obtaining Software Security Patches for your HP Computer System** -- Describes the process for obtaining security patches for HP systems.

Chapter 2, Overview--Communicator Summary

This chapter provides a summary of information contained in this manual. It also provides information about obtaining MPE patches from the HP Electronic Support Center.

Chapter 3, Growth Solution

This chapter provides technical articles about the Growth Solution, including the following topics:

- **Large Files overview** -- describes the changes to the MPE/iX operating systems as a result of the introduction of Large File functionality.
- **AIFs for Large Files** -- describes the four modified intrinsics for handling Large Files.
- **AIFs for Multiple Job Queues** -- describes the enhanced functionality for obtaining queue information about multiple job queues.
- **Increase in Number of TCP Connections** -- describes an enhancement to allow 20000 connections.
- **Dispatcher Multiple Runqueues** -- describes the changes to the Dispatcher module including improved system performance.
- **Increased System Limits** - an overview -- provides a list of the increased system limits.
- **Support for 511 Disks on MPE** -- describes how to configure more than 255 disks.

- Increased Number of DSTs -- describes the increased number of data segments.
- User-defined Job Queues and ALTJOB HIPRI Enhancements -- describes the enhancement to 255 user-defined job queues possible on Release 6.5.
- NEWCI Command (CI Elimination) -- describes the functionality and implementation of the NEWCI command.
- Compatibility Considerations for COBOL and C -- describes the issues for COBOL and C as they relate to Large Files.
- C/iX Supports 64-bit Integers -- describes enhancement to C/iX compiler to provide full support for a 64-bit integer data type.
- FTP Support for Large Files

Chapter 4, DCM

This chapter provides technical articles about the following topics:

- Enterprise Management Solution -- describes a way to manage distributed IT environments using Legato Backup Solution.
- Support Tools Manager (STM) -- describes a method for handling complex data centers replacing the Sysdiag diagnostics system.
- Enhanced Message Source Template for HP 3000 Agents in ITO -- describes a new set of source templates installed on the Management Server.
- HP Secure Web Console -- describes a new HP product which provides secure console access from web browsers.
- TurboSTORE Support for DLT -- describes the capability for fast search on DLT drives.
- DLT4000/DLT7000 Differential Tape Support on MPE/iX -- describes the capabilities of the DLT4000 and DLT 7000 differential tape drives.
- DLT8000 Differential Tape Drives -- describes the capabilities of the DLT8000 differential tape drive.
- TERMDSM: User Interface Change under System Diagnostics -- describes the features of the new diagnostic interface named Support Tools Manager (STM).
- HP SureStore E Disk Array 12H on the HP 3000 -- describes the new HP disk array supported on the HP 3000.
- DVD Support on the HP 3000 -- provides information about replacing older CD-ROM drives with the DVD drive, while still permitting the use of CD-ROM disks.
- DLT 7000 and Legato Storage Node Usage--provides listing of supported Legato Storage Node and DLT Libraries.
- HP SureStore E Disk Array XP256 -- describes the high-capacity, high-speed mass storage, with continuous data availability, ease of service, scalability, and connectivity.

Chapter 5, Internet, Interoperability and Databases

- Large Transactions for IMAGE Users -- describes the solutions for large transactions

Communicator Summary

available for IMAGE users

- XM User Checkpoint Priority Control--describes the Transaction Manager process to write out dirty pages to reduce response time.
- Transact/iX Expanded B-Tree and File Open Enhancement--describes the enhancements to Transact/iX to support B-Trees and File Open.
- Apache for MPE/iX -- describes new server software which enables your HP 3000 to function as a full-featured web server.
- Introducing LDAP C-SDK/iX -- describes a set of client APIs in C language for MPE/iX applications to access X.500 network directories.
- MPE/iX Software Developer's Kit (SDK) for Java, Version 1.2 -- describes the enhancements to the new release of JAVA for MPE/iX.
- HP Driver for JDBC -- provides information about the Java API that enables development of Java applications and applets with a wide range of relational databases.

Chapter 6, Announcing a Post-6.5 Release Patch

This chapter provides an article about the IMAGE/SQL Date Mapping enhancement. It describes migrating your ACTINFO files and using the new commands, UPDATE TYPE and SPLIT.

Chapter 7, How to Order

This chapter provides information on how to order MPE/iX subsystem products.

Chapter 8, Product Release History

This chapter adds product information for MPE/iX Release 6.5 and updates the termination dates in the Supported System Release Matrix table.

Chapter 9, Catalog of User Documentation

This chapter provides two types of manual listings:

- A listing of all new or updated manuals at the time of the MPE/iX 6.5 Release.
- A current listing of the MPE/iX manuals grouped by collections.

MPE/iX Patches on HP IT Resource Center

*by Patch Support Team
Commercial Systems Division*

MPE/iX patches for MPE/iX Release 5.5 and later, are available on the IT Resource Center (previously the HP Electronic Support Center) to all customers.

Features and Benefits

The new patch access and delivery system benefits all MPE/iX customers with:

- Improved overall communication between HP and customers.
- Provision of useful and timely information for patch justification and decision making.
- Reduced system downtime for known problems.
- Reduction of the turnaround time for patch availability and delivery.
- Close to 24*7 access time.
- Unification of the MPE/iX and HP-UX patch delivery process.

Electronic access to patch information and delivery of patches provide three basic services:

1. Access to patch information in an automated, timely and accurate manner.
2. Electronic downloading of patch information and binaries.
3. Proactive notification of new patches via email.

Access Method to the HP Electronic Support Center

To serve customers the IT Resource Center provides World Wide Web access for downloading patches.

Access to World Wide Web Server (www)

IT Resource Center is available through the World Wide Web. World Wide Web access is the easiest, fastest, and most popular method of browsing for patch information and downloading patches. It is more reliable, especially for large patches.

- **U.S. Web accessing address:**
<http://us-support.external.hp.com>
- **European Web accessing address:**
<http://europe-support.external.hp.com>

Electronic Digests

If you want to keep yourself up-to-date on the latest development of MPE/iX patches, you can sign up for the daily Security Bulletin and weekly mpeix_patch Bulletin. Once you have subscribed to these two bulletins, you will receive these digests on a periodic basis via electronic mail. HP IT Resource Center will inform you proactively about newly developed security and GR patches. For more information, refer to the instructions on the IT Resource Center website.

Patch Installation Tools

There are two tools available to install MPE/iX reactive patches, Patch/iX and AUTOPAT. HP recommends the use of Patch/iX for reactive patch installation. Patch/iX has many features and checks to ease and improve the installation process, including:

- A sophisticated patch qualification mechanism to ensure the integrity of your system.
- The ability to perform much of the patch installation process while your system is still up and available to users.
- An option to install a patch or set of patches using the HP Stage/iX Subsystem, which allows the application of a patch to be performed without tapes. For more information on Stage/iX, refer to the *System Software Maintenance Manual* for your release. Patch/iX instructions are also available on the following website:

<http://www.docs.hp.com/mpeix/>

You should use AUTOPAT only if you are familiar with its use, and have a good understanding of MPE/iX patch management.

Patch/iX Installation Document Retrieval

These are the steps for retrieving documents using Patch/iX.

1. Access the HP IT Resource Center WEB site (previously the Electronic Support Center) using the appropriate WEB address for your country.
2. Click on the link, "Individual Patches".
3. Enter: "ITRC User ID" and "password".
4. Click on the link, "MPE/iX Patches".
5. Click on the link, "MPE Patch Installation Guide".
6. Click on the link, "Use Patch/iX or AUTOPAT to install the patch".
7. Click on the link, "Patch/iX Instructions".

Patch/iX Version Identification

To ensure you have the latest version of Patch/iX, on your system do the following:

1. :HELLO MANAGER.SYS, INSTALL
2. :PATCHIX VERSION
3. Compare this version number (for example, B.01.02) with the latest version available for your release on the HP IT Resource Center Patch/iX download page. If you are

running an earlier version than is available, you should download and install the newer one from the download page.

AUTOPAT Installation Document Retrieval

AUTOPATINST is the “DOCID” of the document with instructions to assist you in installing one or more patches needed by your MPE/iX system using the AUTOPAT installation tool.

1. Access the HP IT Resource Center WEB site (previously the Electronic Support Center) using the appropriate WEB address for your country.
2. In the Main Menu, Click on the link, “Search Technical Knowledge Base”.
3. Enter: “ITRC User ID” and “password”.
4. In the Technical Knowledge Base Home page from the pull down menu, Click on “Search By DOC ID” (do not Search by Keyword).
5. In the search field, enter “AUTOPATINST”.
6. Click on the “SEARCH” button.

Create a CSLT Prior to Patch Installation

Before starting any patch application activity, you should always back up your system by creating a Custom System Load Tape and a full backup. This will allow you the flexibility of restoring your system to the previous environment. To create a CSLT, do the following:

1. Log on as `MANAGER.SYS`
2. `:SYSGEN`
3. `>TAPE`

Disclaimer

CAUTION Hewlett-Packard is not liable for errors occurring during data transmission through the Internet. HP assumes no responsibility for the use or reliability of its software on equipment that it has not furnished itself. Furthermore, for customers without a current support contract with HP, HP is not responsible for answering any questions in regard to the use of this patch delivery process.

3 Growth Solution

This chapter contains the following articles about the Growth Solution.

- MPE/iX Large Files Overview
- AIFs for Large Files
- AIFs for Multiple Job Queues
- Increased Number of TCP Connections
- Dispatcher Multiple Runqueues
- Increased System Limits
- Support for 511 Disks
- Increased number of DSTs
- User-Defined Job Queues and `ALTJOB HIPRI` Enhancements
- NEWCI Command (CI Elimination)
- Compatibility Considerations for COBOL and C
- C/iX Supports 64-bit Integers
- FTP/iX Support for Large Files

MPE/iX Large Files Overview

*by Craig Fairchild
Commercial Systems Division*

Introduction

This document describes the changes to the MPE/iX operating system as a result of the initial addition of Large File functionality. In this document, Large Files are defined as any file greater than 4GB - 64KB (4,294,901,760 bytes) which was the maximum file size prior to release 6.5. For the sake of brevity, this value is referred to somewhat inaccurately as “4GB” for the rest of this document.

Overview

The addition of Large File support to MPE/iX increases the maximum file size supported by MPE/iX from the former limit of 4GB to 128GB for release 6.5. In the future, additional increases in supported file size can be expected as customer needs demand.

In this first release of Large File support, two types of files are supported for file sizes beyond 4GB. These are ordinary, fixed length record files, and a new type of KSAM file, KSAM64. A phased approach to Large File support has been taken to get Large File functionality to market in the quickest possible time frame. Several customer surveys have indicated that these two file types are the most needed for Large File support.

Large File support of KSAM files requires the addition of yet another KSAM file type. Due to the internal structure of CM KSAM and KSAM XL files, these file types cannot support file sizes beyond 4GB. The implication of this was that in order to meet our customer demand of Large File KSAM support, a new KSAM file type needed to be defined. This new KSAM file type is the KSAM64 file, or file type 7. KSAM64 files have an improved, generalized internal structure that allows them to support file sizes from one record to the maximum object size of 128GB. KSAM64 files are externally identical to KSAMXL files (and support all CM KSAM intrinsic calls as well).

Large File support has been added to MPE/iX as seamlessly as possible. Large File support has been implemented within the OS so that Large Files are available on all HP 3000 models. Large Files are created by simply specifying a file limit which is sufficiently large. This can be done interactively through the `BUILD` or `FILE` commands, or programmatically through the `FOPEN` and `HPFOPEN` intrinsics. Opening an existing Large File is as transparent as opening any other file. All system intrinsics will work on Large Files in the exact same manner as any normal sized file. As a matter of fact, if your program accesses files through MPE/iX intrinsics, it is highly likely that there will be no need to make any changes at all to operate on Large Files.

A large part of the ease of Large File use lies in the fact that access to files in MPE/iX is done through a record-oriented interface. All interfaces to the OS still deal with signed 32-bit integers to indicate record numbers. This means that there are really two factors that serve to limit maximum file sizes, the maximum number of records in a file, and the maximum object size supported on MPE/iX. The following table shows how the record size

of an ordinary fixed-length record file can affect the maximum size of file.

Record Size	Maximum # of Records	Maximum Object Size	Maximum Resulting File Size
1 byte	2,147,483,647	128GB	2GB
40 bytes	2,147,483,647	128GB	80GB
60 bytes	2,147,483,647	128GB	120GB
80bytes	1,717,986,918	128GB	128GB

About this Document

This document is intended to describe all the external functions, features, and changes that have been introduced with Large File support. The document itself is broken into four sections:

- **Large File Compatibility Issues.** This section covers the small number of issues that could arise when accessing Large Files.
- **Changes to the Command Interpreter.** This section describes the enhancements that have been added to several CI commands to support Large Files and KSAM64 files.
- **New and Enhanced Intrinsic.** This section reveals the upgrades to several existing intrinsic for Large File support, as well as introducing a new set of intrinsic to aid in user-mapped file access.
- **Miscellaneous.** This section discusses the impacts of Large Files on various subsystems and utilities, and covers other miscellaneous Large File topics.

Summary of Large File Compatibility Issues

This section covers the compatibility issues that users need to be aware of when dealing with Large Files. In most cases applications and users are unaware that the file that they are operating on is a Large File.

Even though Large Files are highly transparent, there are some rare situations that can produce results that are either incompatible or in error when dealing with Large Files instead of normally-sized files. The following sections highlight and describe the issues that can affect access to Large Files.

User-mapped Access to Files

MPE/iX currently supports two types of user-mapped access to files; short-mapped access and long-mapped access. Short-mapped access provides a 32-bit pointer for a user application to directly access a file. There are several limitations to short-mapped access. For example, a file must be 4MB or less and the total of all short-mapped access files for a process must be less than 6MB. Long-mapped access provides a 64-bit pointer to access a file of any size, up to the previous maximum file size of 4GB.

Problem Description

Due to the way that existing MPE/iX compilers emit instructions to access memory, problems can result when accessing Large Files with a Large File pointer. Long pointers on PA-RISC hardware are defined to have two separate parts; a space ID, and an offset within that space. This also corresponds to the way virtual memory is defined in PA-RISC -- a 64-bit address broken down into the high 32-bits defining the space of the address, and the lower 32-bits defining the offset within that space.

When building instructions to access memory, the space ID must first be loaded into a space register. The offset portion of the address is then loaded into a general register and the space register/general register pair is used to de-reference virtual memory. When adding to the pointer, the general register is incremented and then used in conjunction with the space register that was previously loaded. When the maximum file size supported on the system was 4GB (the size of each space), all valid file offsets were found within a single space, and there was no need to change the space ID.

The advent of Large Files makes it possible for a single file to cross multiple space ID boundaries. This presents a problem for existing code which uses long-mapped access. When the general register holding the offset is incremented to the point that it overflows, logically it should reference the first word in the next space. However, because the space ID is held in a separate space register, the overflow of the general register is lost. The net effect is that access loops back to the beginning of the current space rather than moving on to the next space as intended.

Compatibility Issue

In order to prevent applications from inadvertently stumbling across this data access problem, normal long-mapped access to Large Files has been made illegal. Any attempt to open a Large File for long-mapped access (HPFOPEN option 21) results in an error. In order to open a Large File for user-mapped access, a new large-mapped option must be passed (HPFOPEN option 87). This new large-mapped option works on both Large Files as well as normal-sized files. In programs that wish to continue to access Large Files through user-mapped access, all instances of HPFOPEN calls for long-mapped access (option 21) need to be changed to large-mapped access (option 87). The pointer that is returned is still a 64-bit pointer which is equivalent in every way to the pointer type returned by the long-mapped option.

By specifying this new large-mapped access option, the option assumes that the caller is indicating their awareness of cross space ID problem and ensures that the program logic correctly handles any cross space ID buffer.

Workaround

The workaround to this problem is actually fairly simple. Any access that would normally span a space ID boundary needs to be broken into two separate accesses, one for each side of the space ID boundary. This can be done within the calling application directly, or by calling the new intrinsics described in section 3 of this document. In order to open a Large File for user-mapped access, the new HOP_OPTION_LARGE_MAPPED (87) option must be passed instead of the long-mapped option.

FLABELINFO

Two current items of the FLABELINFO intrinsic are unable to return information on some Large Files. These items are described in the following table:

Item #	Description	New Item #	Description
28	This item returns the size of the file limit of a file in bytes. It returns the value in an unsigned 32-bit integer. For any Large File, an error is returned if this item is called since the true file limit could not be expressed in a 32-bit value.	62	This item returns the number of bytes in the file limit, but instead returns the value in a signed 64-bit integer. This item can be used for all files regardless of whether or not they are Large Files.
58	This item was recently added to MPE/iX to return the number of sectors of disk space currently allocated to a file. This value is returned in a 32-bit signed integer. While this is adequate for all files supported as of 6.5, future release may support file sizes which will exceed the capacity of this item (512 GB).	63	This item forms a replacement for item 58. It also returns the number of sectors allocated to a file, but returns this information in a signed 64-bit integer. Because of this increased size, this item can reliably be used to get this disk space information for any sized file.

32-bit Integer Overflow

There are a number of FFILEINFO, FGETINFO and FLABELINFO items that return information on the maximum file size of a file, or the current size of a file. While each of these items (with the exception noted above) works without change on Large Files, it is possible that programs that perform arithmetic on the values returned from these functions could experience integer overflow traps when dealing with very large files. An example of such an operation is multiplying the number of records in a file by the record size to get an estimation of the total size of a file. With normal-sized files, the result of this multiplication is guaranteed to never overflow an unsigned 32-bit integer. However, with Large Files, the result of this multiplication could be a value up to 128 GB, which could be much larger than a value that can be represented in 32-bits. In such cases, the result of the multiplication would need to be assigned to a 64-bit integer.

Intrinsic	Item # or Parameter	Item Name	Description
FFILEINFO	4	FF_REC_SIZE	A 16-bit signed integer which returns the record size associated with the specified file

Intrinsic	Item # or Parameter	Item Name	Description
FFILEINFO	9	FF_RECORD_PTR	A 32-bit signed integer which returns the record number of the current logical data pointer associated with the specified file number
FFILEINFO	10	FF_EOF_PTR	A 32-bit signed integer which holds the current EOF record number of the specified file
FFILEINFO	11	FF_FILE_LIMIT	A 32-bit signed integer that returns the record number of the file limit of the file (the maximum size that the file may grow to)
FFILEINFO	14	FF_BLK_SIZE	A 16-bit signed integer that returns the block size associated with the specified file
FFILEINFO	67	FF_HPE_REC_SIZE	A 32-bit signed integer that returns the record size associated with the given file
FFILEINFO	68	FF_HPE_BLK_SIZE	A 32-bit signed integer that returns the record size associated with the given file
FLABELINFO	12	FLI_FILE_LIMIT	A 32-bit signed integer that returns the maximum number of records in the specified file
FLABELINFO	19	FLI_RECORD_END_OF_FILE	A 32-bit signed integer that returns the current data pointer record number associated with the given file number
FGETINFO	N/A	<i>lrecptr</i>	A 32-bit signed integer that returns the current data pointer record number associated with the given file number
FGETINFO	N/A	<i>eof</i>	A 32-bit signed integer that returns the current EOF record number
FGETINFO	N/A	<i>filelimit</i>	A 32-bit signed integer that returns the maximum number of records in the specified file

If any of the items listed above are used in mathematical computations, especially when a record number item is multiplied by a record or block size, care must be taken to determine whether or not it is possible to overflow a 32-bit integer.

AIF:OS Impact

The Architected Interfaces for operating system information (AIF:OS) are a set of privileged routines that provide a supported means of accessing and altering internal system data structures. Because of the need to enhance these data structures to provide Large File support, some corresponding changes to a few AIF:OS routines are required.

Each of the affected AIF:OS items follows a similar form. A 32-bit value is currently being returned, which cannot express the capacity of a Large File. In each case a new item number has been defined which works with a 64-bit value and can be used safely for either large or normal sized files.

The following table describes each of the AIF:OS items that return errors for Large Files, as well as the new items that have been added to supplant the older items:

AIF Function	Item #	New Item #	Description
AIFFILELGET	4033	4101	This item returns the byte offset of the current data pointer associated with the file number that was passed to AIFFILELGET. Item 4033 returns a 32-bit unsigned integer and results in an error when called against a Large File. The 4101 item returns a 64-bit integer and works for all file sizes.
AIFFILELPUT	4033	4101	This item allows a caller to set the data pointer associated with a file number to a new byte offset within a file. Item 4033 allows a caller to pass in a 32-bit unsigned integer to specify the new data pointer value. Given its size, this item can never set the data pointer past the first 4GB of a Large File. The new 4101 item number allows the caller to specify a 64-bit integer to indicate the new data pointer value. The new 4101 item number works for all files, providing the value specified is within the file limit of the specified file.
AIFFILEGGET	5017	5101	These items return the byte offset of the End Of File pointer for a given file. Item 5017 returns a 32-bit unsigned integer and generates an error when called against a Large File. Item 5101 returns a 64-bit integer which can be safely be returned for both large and normal sized files.

AIF Function	Item #	New Item #	Description
AIFFILEGPUT	5017	5101	<p>These items allow the EOF of a file to be set to a specified byte offset. Item 5017 allows the caller to specify a 32-bit unsigned integer for the new EOF offset. Due to the size of this value, it is not possible to use this item to set an EOF beyond 4GB in any file. Item 5101 allows the caller to specify a 64-bit integer which can be used to set the EOF to any offset in any file, providing that the EOF offset is within the file limit of the specified file.</p>
AIFFILEGGET	5018	5102	<p>These items return the byte offset of the file limit of a file, or the maximum size that the file can possibly grow to. Item 5018 returns this value in a 32-bit unsigned integer and returns an error if called against a Large File. Item 5102 returns this information in a 64-bit integer and can be safely used with all file sizes.</p>
AIFFILEGGET	5048	5103	<p>These items return the number of sectors of disk space that currently are allocated to a given file. Item number 5048 returns this value in a signed 32-bit integer and returns an error for a Large File, if that Large File has more than 512GB of disk space allocated to it. Item 5103 returns this sector count in a 64-bit integer and reliably works on files of any size.</p>

OPTIONS Enhancement

As a side effect of adding support for the KSAM64 file type, a new file type value, 7, has been allocated for KSAM64 files. The `foptions` structure is a 16-bit value made up of various bit fields. Several intrinsics accept or return `foptions` including `FOPEN`, `FFILEINFO`, `FLABELINFO`, and `FGETINFO`. The definitions of the `foptions` are:

0:1	Reserved
1:1	Record format extension bit
2:3	File Type
5:1	File Equations disallowed
6:1	Labeled tape
7:1	Carriage control
8:2	Record format
10:3	Designator type
13:1	Binary/ASCII
14:2	File domain

The addition of KSAM64 support in no way alters the format of the `foptions` structure. Rather, a previously undefined value in the file type field, 7, is now defined. The list of defined file types which can be returned in the `foptions` is:

- 0 Ordinary file type
- 1 CM KSAM
- 2 RIO
- 3 KSAMXL
- 4 CIR
- 5 SPOOL
- 6 MSG
- 7 KSAM64

In general, the definition of this previously undefined value will not cause any compatibility issues for programs. The only time that issues could arise is if code performed some type of validity check on the file type field of the `foptions`. In the case that user code is checking and rejecting a value of 7 in the file type field, the workaround would be to eliminate this check, or add 7 to the list of valid values.

Command Interpreter Changes

Although the impact of Large Files is negligible on most CI commands, there are several minor enhancements to a few commands. All CI commands work on Large Files in the exact same manner as they do for normal sized files. `RENAME` still renames Large Files, `PURGE` still purges them, and so on. The commands documented in this section have been enhanced to support some Large File features.

BUILD Command Changes

With the first release of Large File support, only two file types are being supported for Large File sizes; ordinary, fixed length record files, and KSAM files. Because of inherent limitations in the internal format of KSAMXL and CM KSAM files, it is necessary to introduce a new type of KSAM file, KSAM64 files, in order to go beyond a 4GB sized KSAM file. The internal format of the KSAM64 index area is structured so that file sizes from 1 record all the way to 1TB can be easily supported.

While unfortunately, this means that MPE/iX supports three different varieties of KSAM files (CM KSAM, KSAMXL and KSAM64), the good news is that all three of these file types are functionally equivalent and can be used interchangeably by programs. The only external difference between KSAMXL and KSAM64 files is that a different file type must be used to create them. Other than this, all other KSAM options are specified identically.

The only change to the BUILD command is the addition of the KSAM64 file type keyword. This keyword can be included in the BUILD command as a replacement for the KSAMXL keyword. KSAM64 files may be of any file size, but are the only KSAM file type that can be larger than 4GB. The new syntax for the BUILD command is:

```
BUILD filereference

                                [F]
[;REC=[[recsize][,[blockfactor][,[U][,BINARY]]]]]
                                [V][,ASCII ]
                                [B]

[ {;CCTL  } ]
  {;NOCCTL}
[;TEMP]
[;DEV=[[dsdevice]#[device]]
[;CODE=filecode]
[;DISC=[numrec][,[numextents][,[initialloc]]]]

[ {;RIO  } ]
  {;NORIO}

  {;STD  }
  {;MSG  }
[ {;CIR  } ]
  {;KSAMXL}
  {;SPOOL }
  {;KSAM64}

[;ULABEL=numlabels]
[;KEY={^filereference2}]
      {keyinfo      }

[;FIRSTREC=recnum]
[ {;REUSE  } ]
  {;NOREUSE}

[;LANG={langid  } ]
      {langname}

[ {;DEFBLK} ]
  {;OPTMBLK}
```

In order to create a Large File, all that is required is to specify a file limit of a sufficiently large number. For example, to create an ordinary, fixed length record file with 80 byte

records that had a maximum capacity of about 8GB, a file limit of 100,000,000 records would need to be specified. For example, the following command could be used:

```
BUILD TESTFILE;REC=-80,,F,ASCII;DISC=100000000
```

Specifying a file limit that is greater than 4GB for any type of file other than ordinary fixed length record files and KSAM64 files will result in an error.

FILE Command Changes

The FILE and BUILD command syntaxes always mirror each other. It follows, then, that the new KSAM64 file type keyword is also being added to the FILE command, in the exact same way as the BUILD command. The new FILE syntax is included below:

```
                =*formaldesignator
                =$NULL
                =$NEWPASS
FILE formaldesignator =$OLDPASS
                =$STDIN
                =$STDINX
                =$STDLIST
                =filereference[:nodespec]
                                [,filedomain]

[;DEV=[[envname]#][device][,outpri][,numcopies]
[;VTERM]
[;ENV=envfile[:nodespec]]
[;option]
[;access]
[;disposition]
[{:DEFBLK}]
{:OPTMBLK}
```

option Any valid option for the FILE command.

SYNTAX FOR OPTION

```
                {F}
;REC=[recsize][,[blockfactor][,{U}[,{BINARY}]]]
                {V} {ASCII }

;DEN=[density]
;DISC=[numrec][,numextents][,initialloc]
;CODE=[filecode]

[ {;RIO } ]
{:;NORIO}
{:;STD}
{:;KSAMXL}
{:;SPOOL}
{:;KSAM64}

[ {;MSG} ]
{:;CIR}

[;ULABEL=numlabels]

[;KEY={^filereference2}]
{:keyinfo}

[;FIRSTREC=recnum]

[ {;REUSE} ]
{:;NOREUSE}

[;LANG= { langid } ]
{: langname}
```

LISTF[ILE | TEMP] Command Changes

The LISTF, LISTFILE, and LISTFTEMP commands all can be used to display information about Large Files. All of the various LISTF and LISTFILE formats continue to work with normal sized files as well as Large Files. File type information about KSAM64 files is displayed in the LISTF family of commands as a lowercase 'k'.

In order to display information about Large Files, two new file format options have been introduced. The LISTF, LISTFILE, and LISTFTEMP commands now support options 10 (SUMMARYWIDE) and 11 (DISCWIDE). File format option 10 is the Large File equivalent of file option "1", containing all the same information, but in a new format that allows for the expression of larger values. Similarly, the format option 11 is the Large File equivalent of file option "2". Again, the same basic information is displayed, but the format has been changed to allow for a greater range of values to be displayed.

The syntax of the LISTF, LISTFILE, and LISTFTEMP commands has not been changed with the addition of the new 10 and 11 format options. The same syntax can be used with the new values of 10 and 11 as the format. The output of the new formats varies depending on whether or not MPE/iX syntax names are being displayed or HFS syntax names are being displayed.

The format for each of the MPE/iX and HFS syntax varieties of the new formats 10 and 11 can be seen from the examples shown below:

```
:listf @,10
```

```
ACCOUNT=  SYS          GROUP=  EXAMPLE
Name      Access FCode RecSiz Type          EOF      File Limit
-----  ERWS  -----  -----  -----  -----  -----
BIGFILE           1024 FA          0      1023456789
ICE              E      NMPRG   256 FB          832      832
XKSM64           80 FAK          1      1023
XKSMXL           80 FAK          1      1023
XRAND            RW      80 FA          157      157
XRAND2           W       252 VA          396      49
YRAND            NMOBJ   256 FB          22      4000
```

```
:listfile ./@,10
```

```
PATH= /SYS/EXAMPLE/
```

```
Access FCode RecSiz Type          EOF      File Limit Name
ERWS  -----  -----  -----  -----  -----  -----
      1024 FA          0      1023456789 BIGFILE
E      NMPRG   256 FB          832      832 ICE
      80 FAK          1      1023 XKSM64
      80 FAK          1      1023 XKSMXL
RW     80 FA          157      157 XRAND
W      252 VA          396      49 XRAND2
      NMOBJ   256 FB          22      4000 YRAND
```

```
:listf @,11
ACCOUNT=  SYS          GROUP=  EXAMPLE
```

Name	Access	FCode	RecSiz	Type	EOF	File Limit	Disk Usage	Exts
-----	ERWS	-----	-----	-----	-----	-----	-----KB	-----
BIGFILE			1024	FA	0	1023456789	0	0
ICE	E	NMPRG	256	FB	832	832	208	1
XKSM64			80	FAk	1	1023	64	1
XKSMXL			80	FAK	1	1023	52	1
XRAND	RW		80	FA	157	157	16	1
XRAND2	W		252	VA	396	49	16	1
YRAND		NMOBJ	256	FB	22	4000	8	1

```
:listfile ./@,11
PATH= /SYS/EXAMPLE/
```

Access	FCode	RecSiz	Type	EOF	File Limit	Disk Usage	Exts	Name
ERWS	-----	-----	-----	-----	-----	-----KB	-----	-----
		1024	FA	0	1023456789	0	0	BIGFILE
E	NMPRG	256	FB	832	832	208	1	ICE
		80	FAk	1	1023	64	1	XKSM64
		80	FAK	1	1023	52	1	XKSMXL
RW		80	FA	157	157	16	1	XRAND
W		252	VA	396	49	16	1	XRAND2
	NMOBJ	256	FB	22	4000	8	1	YRAND

The display for the format options 10 and 11 is the same for each of the LISTF, LISTFILE, and LISTFTEMP commands. One of the unique additions of these two new options is the “ERWS” column. This column indicates whether or not the file is currently opened, and if so, for what type of access. An “E” in the E column indicates that the file is opened for Exclusive access, the “R” in the R column indicates that the file is opened for Read access, the “W” in the W column indicates that the file is opened for Write access, and the “S” in the S column indicates that the file is being Stored by the STORE or TurboSTORE utility. The column may have one or more of these characters specified, or none, depending on how the file is currently accessed.

Intrinsics Changes

As part of the addition of Large File support to MPE/iX several intrinsics changes have been introduced. This encompasses both the modification of existing intrinsics, as well as the introduction of a new set of intrinsics. This chapter describes all of the intrinsics changes that have been implemented as a part of Large File support.

Existing Intrinsics

Relatively few changes have been necessary to the existing file system intrinsics in order to implement Large Files. The majority of these changes are in support of adding 64-bit return values for file byte offsets.

HPFOPEN Changes

As described in the section “User-Mapped Access to Files”, when performing user-mapped access to Large Files, care must be taken to not attempt to dereference a buffer that

crosses a SID boundary. In order to ensure that programs are aware of this caveat, traditional attempts to open a Large File for user-mapped access will fail. In order to open a Large File for user-mapped access, a new HOP_OPTION_LARGE_MAPPED option (item number 87) must be used, as described in the following table:

Item Name	Item Number	Item Size	Description
HOP_OPTION_LARGE_MAPPED	87	@64	This item returns a 64-bit pointer value which points to the start of the file data of the specified file. This item works for any sized file, but is the only user-mapped access option that works for Large Files (files whose file limit is greater than 4GB). Note that not all file types allow user-mapped access. If this item is requested against a file type that does not support user-mapped access, the HPFOPEN request will result in a failure.

FLABELINFO Changes

Three new items will be supported by the FLABELINFO intrinsic. These new options are described in the table below:

Item Name	Item Number	Item Size	Description
FLI_FILE_LIMIT_BYTES_64	62	8 bytes	This item returns a 64-bit integer which indicates the maximum file size of the specified file in bytes. This is the 64-bit equivalent of FLABELINFO item 28.
FLI_COMPUTED_SECTORS_64	63	8 bytes	This item returns a 64-bit integer which holds the number of 256 byte sectors of disk space that are currently allocated to the specified file. This item is the 64-bit equivalent to the FLABELINFO item 58.
FLI_FILE_LIMIT_IS_LARGE	64	4 bytes	This item returns a 32-bit integer which can be used as a simple test to see if the file limit of a file is greater than 4GB. This item will return two values: 0 = The file limit is less than 4GB 1 = The file limit is greater than 4GB

FFILEINFO Changes

Two new items have been added to the `FFILEINFO` intrinsic. These items are described in the table below:

Item Name	Item Number	Item Size	Description
<code>FF_FILE_LIMIT_BYTES</code>	109	8 bytes	This item returns a 64-bit integer indicating the maximum size of the file, or the file limit, in bytes.
<code>FF_FILE_LIMIT_IS_LARGE</code>	110	4 bytes	This item returns a 32-bit integer which can be used to easily tell if the specified file has a file limit greater than 4GB. This item returns two values: 0 = The file limit is less than or equal to 4GB 1 = The file limit is greater than 4GB.

New Intrinsics

The following new intrinsics are being added to aid in the support of user-mapped file access. These intrinsics can be used for pointer manipulation and data movement on any user-mapped file; short-mapped, long-mapped, or large-mapped. They are especially recommended for use with the large-mapped `HPFOPEN` option.

HPFADDTOP_POINTER

NM callable only.

This routine can be used to perform arithmetic on a 64-bit pointer value. Byte offsets can be added to or subtracted from a pointer by specifying either a positive or negative offset value.

Syntax

```

                                @64      I64      @64      I32
HPFADDTOP_POINTER( base_ptr, offset, return_ptr, status );
```

Parameters

- base_ptr* **64-bit pointer by reference (required)**
The *base_ptr* can be a 64-bit pointer to an object of any type.
- offset* **64-bit signed integer by reference (required)**
The offset can be any positive or negative value. Specifying a positive value moves the *return_ptr* forward from the previous *base_ptr*, while a negative value moves the *return_ptr* backward from the *base_ptr*.
- return_ptr* **64-bit pointer by reference (required)**
The *return_ptr* is an output parameter that has the new pointer value returned to it. It can be a 64-bit pointer to an object of any type.
- status* **32-bit signed integer by reference (optional)**
Returns the status of the `HPFADDTOP_POINTER` call. If no errors or warnings

are encountered, *status* returns 32 bits of zero. If errors or warnings are encountered, *status* is interpreted as two 16-bit fields. Bits (0:16) comprise *status.info*. A negative value indicates an error condition, and a positive value indicates a warning condition. Bits (16:32) comprise *status.subsys*. The value represents the subsystem that set the status information.

NOTE No attempt is made to verify that the pointer value returned is a legitimate pointer to a valid object. Any invalid pointers are detected and generate errors when the pointers are dereferenced.

HPFMOVEDATA

NM callable only.

This routine can be used to efficiently move data from a source buffer to a target buffer.

Syntax

```

                I64          @64          @64          I32
HPFMOVEDATA ( count, source_ptr, target_ptr, status )

```

Parameters

count **64-bit signed integer by reference (required)**
The *count* parameter allows the caller to specify the number of bytes to move from the source buffer to the target buffer.

source_ptr **64-bit pointer by value (required)**
The *source_ptr* can be a 64-bit pointer to any valid object that the calling process has access to. The buffer may be in the caller's stack, heap, or obtained by opening a file with user-mapped access.

target_ptr **64-bit pointer by value (required)**
The *target_ptr* can be a 64-bit pointer to any valid object that the calling process has access to. The buffer may be in the caller's stack, heap, or obtained by opening a file with user-mapped access.

status **32-bit signed integer by reference (optional)**
Returns the status of the HPFMOVEDATA call. If no errors or warnings are encountered, *status* returns 32 bits of zero. If errors or warnings are encountered, *status* is interpreted as two 16-bit fields. Bits (0:16) comprise *status.info*. A negative value indicates an error condition, and a positive value indicates a warning condition. Bits (16:32) comprise *status.subsys*. The value represents the subsystem that set the status information.

NOTE When calling the HPFMOVEDATA intrinsic, it is important to ensure that the source and target buffers are not overlapping. The results of a HPFMOVEDATA call when source and target buffers are overlapping are undefined. If source and target buffers are overlapping, the HPFMOVEDATALTOR or HPFMOVEDATARTOL intrinsics should be used.

HPFMOVEDATALTOR

NM callable only.

This routine can be used to efficiently move data from a source buffer to a target buffer. If the source and target buffers were viewed horizontally, like a line of text, the data movement is performed by starting at leftmost position of the source buffer (to the leftmost position of the target buffer) and proceeding to the rightmost.

Syntax

```
HPFMOVEDATALTOR ( I64count, @64source_ptr, @64target_ptr, I32status )
```

Parameters

<i>count</i>	64-bit signed integer by reference (required) The <i>count</i> parameter allows the caller to specify the number of bytes to move from the source buffer to the target buffer.
<i>source_ptr</i>	64-bit pointer by value (required) The <i>source_ptr</i> can be a 64-bit pointer to any valid object that the calling process has access to. The buffer may be in the caller's stack, heap, or obtained by opening a file with user-mapped access.
<i>target_ptr</i>	64-bit pointer by value (required) The <i>target_ptr</i> can be a 64-bit pointer to any valid object that the calling process has access to. The buffer may be in the caller's stack, heap, or obtained by opening a file with user-mapped access.
<i>status</i>	32-bit signed integer by reference (optional) Returns the status of the HPFMOVEDATALTOR call. If no errors or warnings are encountered, <i>status</i> returns 32 bits of zero. If errors or warnings are encountered, <i>status</i> is interpreted as two 16-bit fields. Bits (0:16) comprise <i>status.info</i> . A negative value indicates an error condition, and a positive value indicates a warning condition. Bits (16:32) comprise <i>status.subsys</i> . The value represents the subsystem that set the status information.

NOTE This intrinsic is especially useful when the source and target buffers are overlapping. The HPFMOVEDATALTOR intrinsic is typically used when the target buffer's address is to the left (smaller) of the source buffer's address. Moving the data from the left to the right ensures that the data in the source buffer is copied to the target buffer before it is overwritten itself.

HPFMOVEDATARTOL

NM callable only.

This routine can be used to efficiently move data from a source buffer to a target buffer. If the source and target buffers were viewed horizontally, like a line of text, the data movement is performed by starting at the rightmost position of the source buffer (to the rightmost position of the target buffer) and proceeding to the leftmost.

Syntax

```

                                I64          @64          @64          I32
HPFMOVEDATARTOL ( count, source_ptr, target_ptr, status )

```

Parameters

<i>count</i>	64-bit signed integer by reference (required) The <i>count</i> parameter allows the caller to specify the number of bytes to move from the source buffer to the target buffer.
<i>source_ptr</i>	64-bit pointer by value (required) The <i>source_ptr</i> can be a 64-bit pointer to any valid object that the calling process has access to. The buffer may be in the caller's stack, heap, or obtained by opening a file with user-mapped access.
<i>target_ptr</i>	64-bit pointer by value (required) The <i>target_ptr</i> can be a 64-bit pointer to any valid object that the calling process has access to. The buffer may be in the caller's stack, heap, or obtained by opening a file with user-mapped access.
<i>status</i>	32-bit signed integer by reference (optional) Returns the <i>status</i> of the HPFMOVEDATARTOL call. If no errors or warnings are encountered, <i>status</i> returns 32 bits of zero. If errors or warnings are encountered, <i>status</i> is interpreted as two 16-bit fields. Bits (0:16) comprise <i>status.info</i> . A negative value indicates an error condition, and a positive value indicates a warning condition. Bits (16:32) comprise <i>status.subsys</i> . The value represents the subsystem that set the status information.

NOTE This intrinsic is especially useful when the source and target buffers are overlapping. The HPFMOVEDATARTOL intrinsic is typically used when the target buffer's address is to the right (larger) of the source buffer's address. Moving the data from the right to the left ensures that the data in the source buffer is copied to the target buffer before it is overwritten itself.

HPFFILLDATA

NM callable only.

This routine can be used to efficiently initialize a buffer with a specified character value.

Syntax

```
HPFFILLDATA ( I64 count, @64 buffer_ptr, CV fill_char, I32 status )
```

Parameters

<i>count</i>	64-bit signed integer by reference (required) A positive count of the number of bytes in the buffer indicated by the <i>buffer_ptr</i> parameter that should be initialized.
<i>buffer_ptr</i>	64-bit pointer by value (required) A pointer to the buffer that should be initialized. The <i>buffer_ptr</i> may point to any valid object in your stack, heap, or a file that has been opened with user-mapped access.
<i>fill_char</i>	Character value by value (required) The character value that should be used to initialize the specified buffer. Any value in the range of 0 through 255 can be specified, including all printable and non-printable ASCII characters.
<i>status</i>	32-bit signed integer by reference (optional) Returns the status of the HPFFILLDATA call. If no errors or warnings are encountered, <i>status</i> returns 32 bits of zero. If errors or warnings are encountered, <i>status</i> is interpreted as two 16-bit fields. Bits (0:16) comprise <i>status.info</i> . A negative value indicates an error condition, and a positive value indicates a warning condition. Bits (16:32) comprise <i>status.subsys</i> . The value represents the subsystem that set the status information.

Miscellaneous

This section covers several issues that are related to Large Files that may be of interest to the reader.

STORE/RESTORE Impact

The STORE and RESTORE utilities have been modified to handle Large Files. This enhancement was accomplished in such a fashion that STORE tapes with Large Files are completely compatible with previous versions of the operating system. A STORE tape created on release 6.5 can be taken to any other MPE/iX system and be successfully restored. On systems prior to 6.5, any Large File on the tape is not recognized as a valid file and will not be restored onto the system. Although Large Files are skipped over in this scheme, all other files will be able to be restored.

Any attempt to store a Large File using the ;TRANSPORT STORE option results in an error at store time.

In addition to the STORE/RESTORE and TurboSTORE products all supporting Large Files, most, if not all, 3rd party backup products have also been enhanced to support Large Files.

If you rely on 3rd party backup tools, please refer to your vendor for information regarding the proper updates and procedures for using these tools on release 6.5.

Backdating Procedures

If for some reason it is necessary to “backdate” your system from the 6.5 release to a prior MPE/iX release, very little extra work is required due to Large Files. Before attempting the backdate operation, the 6.5 system must be shut down completely with the `CNTRL-A SHUTDOWN` command. You must also wait for the final system shutdown message to appear before resetting the system and beginning the `UPDATE` command on the earlier MPE/iX release.

Any Large Files that are in the directory structure of the system are unusable once the system is backdated to a release earlier than 6.5. The Large Files can be purged from the older system via the `:PURGELINK` command. Any other attempt to open or manipulate a Large File results in an error. Once the system is again updated to release 6.5 or beyond, the Large Files are once again accessible.

Utility and Subsystem Support

Nearly all utilities and subsystems support Large Files without change. This runs the spectrum from `FCOPY` to `SORT`. The only word of caution is that not all utilities and subsystems support file sizes that are less than 4GB. This includes such things as `EDITOR` or `HPEDIT`. The general rule of thumb is if it doesn't work on a 3GB file, it won't work on a Large File.

One exception of note is the `DSCOPY` subsystem. `DSCOPY` has not been enhanced to be able to transfer Large Files, or the new `KSAM64` files. The `FTP` utility has been enhanced to do network transfers of these files, and is the preferred utility for network file transfer.

Impact on IMAGE Database

The introduction of Large Files has had no impact on the `IMAGE/SQL` product. Large database support is still accomplished by use of Jumbo Datasets. Therefore there are no compatibility issues to be concerned with in regards to changes in your `IMAGE/SQL` databases.

HP SRC

Version A.01.29 of `HP SRC` does not support `KSAM64` files. If an attempt is made to check in a `KSAM64` file, error 387 will be reported.

JAVA

Java programs will be able to access large MPE/iX files by using Java Native Interface (JNI) to call MPE/iX intrinsics, but will not be able to access Large Files via the `java.io` APIs. The `java.io` APIs treat all files as bytestream files and access these files through the POSIX C language APIs. The first release of Large Files for MPE/iX does not include support for bytestream files larger than 2GB in size; therefore, there can be no Large Files that are also bytestream files in this release.

HPSEARCH/iX

`HPSearch/iX` has been changed to provide for searching of Large Files. Additionally, it has

been enhanced to search KSAM64 and KSAMXL files.

HPGLANCEPLUS/iX

HPGlancePlus/iX has been changed to display file information about Large Files. Also, KSAM64 files are included in the File Display on the Process Screen and will show a type of KSM64.

HPBROWSE/iX

HPBrowse/iX supports viewing Large Files and KSAM64 Files.

HPEDIT/iX

HPEdit/iX does not support files of 2,000,000 records or more, and so, does not support Large Files.

COBOL/iX

COBOL II/iX supports Large Files. No changes are necessary to existing COBOL II/iX programs to read, write, and update Large Files. KSAM64 files are supported by the Indexed I-O module. When you open a file with indexed organization, that file is permitted to be an existing KSAM64 file. When COBOL II/iX creates a file with indexed organization, by default it will create a KSAM XL file. However, you may override the default by issuing a FILE command specifying KSAM64 as the file type.

COBOL II/CM

COBOL II/CM does not support KSAM64 files. An attempt to open or create a KSAM64 file with a CM COBOL program will result in file status 39 and, typically, a program abort with COBOL error 648.

C/iX

C/iX does not support Large Files. You must issue your own intrinsic calls if you want to access Large Files from C. Attempting to open a Large File using the C/iX Library function `fopen()` will fail and set `errno` to `ESYSERR (50)`.

AIFs for Large Files

by Rajesh Channabasavaiah
Commercial Systems Division

The Architected Interface Facility (AIF) Operating System product has been enhanced to handle the instances of Managing Large Files (Up to 1TB) with the AIF Intrinsic. AIFs are a software layer between non-operating system software and internals providing controlled access to MPE/iX internal functionality and data structure. For Release 6.5, four AIF Intrinsic have been identified and modified to obtain the required functionality for handling Large Files.

The four AIF Intrinsic are as follows:

- AIFFILEGGET
- AIFFILEGPUT
- AIFFILELGET
- AIFILELPUT

Enhancement for AIFFILEGGET/AIFFILEGPUT

These Intrinsic have been enhanced for handling Large Files with the addition of three item numbers which are as follows:

Item #	Item Type	Description
5101	Longint_type	End of File
5102	Longint_type	File limit
5103	Longint_type	Number of Sectors

Enhancement for AIFFILELGET/AIFILELPUT

These Intrinsic have been enhanced for handling Large Files with the addition of a New Item number which is as follows:

Item #	Item Type	Description
4101	Longint_type	Record pointer offset

Error Messages

Two New Error Messages have been introduced for the AIF:OS product as part of the Large Files project to handle the existing items and the new items for the Intrinsic documented above.

Error #	Description
4101	Set if the designated file is a Large File
4102	Set if the designated file is not a Large File

AIFs for Multiple Job Queues

by *Rajesh Channabasavaiah*
Commercial Systems Division

The Architected Interface Facility (AIF) OS product has been enhanced to return the information about the multiple job queues running on the MPE/iX Version 6.0 and above. The Job Queue information obtained will provide the job queue name, job queue limit, number of jobs executing and also the total number of jobs in the queue.

To obtain the above requirement for the AIF, two AIF intrinsics have been enhanced. The two intrinsics are as follows:

- AIFSYSWIDEGET
- AIFJSGET

Enhancement for AIFSYSWIDEGET

This intrinsic has been enhanced to return the complete job queue information based on the job queue name specified. The user can request the job queue information like the limit, jobs executing and total jobs executing for a particular job queue name. By passing that job queue name as an input via the AIFSYSWIDEGET call through the parameter ITEM_ARRAY, all the required information about that job queue name is returned in the return parameter, RETURN_ARRAY1 .

If the requested job queue name is not a valid name in the existing job queue list, then error number 1014 is returned which indicates that the requested job queue name does not exist in the queue list.

Enhancement for AIFJSGET

This intrinsic has been enhanced to return the individual job queue information, depending on the item number for which the information is required. The job number has to be specified when this intrinsic is called. The individual item numbers for obtaining the information for each item are as follows:

Item Number	Type	Description
1051	CA8	Job Queue Name
1052	I32	Job Queue Limit
1053	I32	Number of Job Queues executing
1054	I32	Total Number of Job Queues executing

Increase in Number of TCP Connections

*By Ganesh HS & Chandrashekar MN
Commercial Systems Division*

Introduction

Prior to the MPE/iX 6.5 release, TCP on MPE/iX allowed a maximum of 10240 connections. HP has enhanced the TCP module to get up to 20000 connections.

Configuration through NMMGR

You need to configure the machine to allow it to accommodate up to 20,000 connections. The NMMGR screen for configuring the maximum connections looks like this -

```
/*-----*
/
NMMGR/3000 (B.06.00)  Transmission Control Protocol (TCP) Config   Data: Y
When Data Flag is "N", press "Save Data" to create the data record.Command:
Path:  NETXPORT.GPROT.TCP

[N]      Checksum Enabled (Y For Yes, N For No)
[15612]  Maximum Number of Connections
[2  ]   Retransmission Interval Lower Bound (Secs)
[180]   Maximum Time to Wait For Remote Response (Sec)
[10 ]   Initial Retransmission Interval (Secs)
[4  ]   Maximum Retransmissions per Packet
[600]   Connection Assurance Interval (Secs)
[3  ]   Maximum Connection Assurance Retransmissions

File:  NMCONFIG.PUB.SYS

                                         Save   Help   Prior
                                         Data   Screen

/*-----*
/
```

The field, "Maximum Number of Connections" can now handle up to 20,000 connections.

To reach this screen, “Transmission Control Protocol (TCP) Config” in NMMGR, do the following:

1. Go to [Open Config]
2. Go to [NS]
3. Go to [Unguided]
4. Go to [Netxport]
5. Go to [Gprot]
6. Go to [TCP]

After the re-configuration, stop the network (via `NETCONTROL`) if it is already started and restart the network. Now TCP is ready to accommodate up to 20,000 TCP connections.

Dispatcher Multiple Runqueues

by Senthil Kumar R
Commercial Systems Division

Introduction

The MPE/iX Operating System is continuously undergoing various enhancements to keep pace with the increasing performance needs of HP 3000 customers. The support for Dispatcher Multiple Runqueues is one such enhancement in MPE/iX Release 6.5 which is aimed at providing improved system performance to customers. Dispatcher is one of the core modules of the MPE/iX OS; it is responsible for allocating CPU time to processes. This article gives a technical overview of the changes to the Dispatcher module to provide improved performance on 10-way and 12-way systems in MPE/iX 6.5.

Technical Overview

Dispatcher maintains the list of ready-to-run processes in the system in a queue called “runqueue”. Earlier versions of MPE/iX used a single runqueue to maintain these processes. This was true even in multiprocessor (MP) systems which have more than one processor. So in an MP system all processors have to contend for this single runqueue while picking up processes for execution. When one processor is accessing the runqueue, other processors wanting to access the queue have to wait until the runqueue is free. This serialization causes severe contention in the system and leads to loss of performance, especially in higher-end multiprocessor systems.

To overcome the performance loss due to runqueue contention, MPE/iX 6.5 supports multiple runqueues in Dispatcher. There can be as many runqueues as the number of processors on the system. The ready processes can be enqueued to any of these runqueues. With multiple runqueues available, each processor can pick ready processes for execution independent of other processors. This implies that all processors can now select in parallel ready processes for execution instead of doing it serially as in the case of a single runqueue. The processors do not waste time waiting on the runqueue and can do more useful work. This results in improved system throughput and performance.

Benefits to Customers

The support for multiple runqueues in Dispatcher brings many advantages for MPE/iX customers. With this feature, the scalability of MPE/iX has been considerably increased. This means that the increase in system performance with the addition of more processors is much more than what it was earlier. This is a key aspect to enable customers to upgrade their HP 3000 systems with a larger number of processors to fulfill their performance needs. Also the system response time for users is expected to be better since process scheduling will be faster now.

In addition to supporting multiple runqueues, there have been other changes in the Dispatcher module to make it more efficient for execution in multiprocessor environment. This has been done by carefully rearchitecting various synchronization mechanisms in the implementation to cut down overheads and focus on delivering faster response time.

Conclusion

In summary, the support of multiple runqueues in MPE/iX 6.5 boosts the performance of the HP 3000 to a great extent and makes it more scalable with respect to number of processors used. It has been enabled for 997 systems with 10 or 12 processors on Release 6.5. The users can also experience better response time from the system.

Increased System Limits - An Overview

*by Ganapati
Commercial Systems Division*

One of the major enhancements of the MPE/iX 6.5 release is increases to several former operating system limits. These enhancements are designed to help large HP 3000 sites continue to grow the workload on their systems.

It is important to note that some of these limits can be achieved only under certain conditions. Further details will be provided in the articles which follow.

Operating System Limits	Previous Limits	New Limits
Main Memory	3.75GB	16GB
Disk devices per system	255	511
Disk devices per device class	255	511
DSTs per system	16383	65535
Ordinary fixed-length record files	4GB	128GB
KSAM Files	4GB	128GB
VT Sessions	2000	2600
TIO Connections	2750	3300
TCP Connections	10240	20000
UDP Connections	4096	10000

Support for 511 Disks on MPE/iX

by N. A. Narayanan
Commercial Systems Division

Introduction

Beginning with MPE/iX Release 6.5, it is possible to configure up to 511 spindles, an increase from the present limit of 255 disks.

Configuring More Than 255 Disks

The existing utilities `volutil` and `mirvutil` are used to configure more than 255 disks. There are no interface changes to these utilities to configure more than 255 disks.

- However, it is not possible to configure more than 255 disks in a volume set. This was the same limit present in 6.0. HP recommends that customers use multiple volume sets when a large number of disks are installed on a single system.
- Also, it is possible to configure 255 mirrored pairs (510 disks) in a mirrored volume set.

User Interface Changes

- Sysgen has been modified to configure up to 511 disks in a device class. Until 6.5, sysgen could configure only up to 255 devices in a device class.
- Utilities like `volutil`, `mirvutil`, `dstat`, `discfree`, etc. can display information related to more than 255 disks now.
- `Mirvutil` can configure up to 255 mirrored pairs (510 disks) in a single mirrored volume set now.
- Apart from the above, there are no changes in the interface of any of the utilities.

Forward/Backward Compatibility Issues

There are no compatibility issues between 6.5 and 6.0 or between 6.5 and 5.0/5.5. A system that has been installed with 6.5 can be back-dated to 6.0 without any problem, as long as it does not have more than 255 spindles configured. The same holds good for back-dating from 6.5 to 5.0/5.5.

Increased Number of DSTs

by *Senthil Kumar R*
Commercial Systems Division

Introduction

DSTs are Data Segments used in MPE/iX for supporting Compatibility Mode (CM) operation. With the increase in capabilities of HP 3000 systems and MPE/iX OS over the years, it became necessary to increase the total number of DSTs supported in MPE/iX. This article gives a technical overview of the changes done in MPE/iX Release 6.5 to increase the number of available DSTs from 16383 to 65535. It also describes the benefits that this change can bring to HP 3000 customers.

Technical Overview

Earlier versions of MPE/iX supported a total of 16383 DSTs. This limit of 16383 DSTs was mainly due to limitations in some internal data structures of the OS that were used to handle DST numbers. These limitations have been adequately addressed in MPE/iX 6.5 to allow handling of larger DST numbers. This has helped in increasing the total number of DSTs from 16383 to 65535.

Prior to MPE/iX 6.5, we had a total of 16383 DSTs which were shared by both OS as well as user applications. With the 6.5 release, we distinguish between DSTs used by the OS and DSTs allocated by user application programs. We now have about 16383 DSTs exclusively available for user applications. The OS has a separate pool of 49152 DSTs for its own internal usage. These two together make a total of 65535 DSTs available in the system. The increase allows more DSTs to be available to satisfy the requirements of the OS as well as user applications.

Benefits to Customers

The increase in number of DSTs brings many advantages for MPE/iX customers. Since DSTs are mainly used inside the OS to manage jobs and sessions, an increase in the number of DSTs available implies that the system can now run more jobs and sessions compared to earlier releases. The increase in number of jobs/sessions can, in turn, help in increasing the number of simultaneous users supported on the system.

With 16383 DSTs now exclusively available for user applications, users can run more CM applications that make use of DSTs at application level.

Since MPE/iX Networking modules are also major consumers of DSTs, the increase in DSTs facilitates capacity improvement in Networking modules. For example, the DST increase has contributed largely to the increased number of Virtual Terminal (VT) connections available from MPE/iX 6.5 onwards.

Conclusion

In summary, the increase in number of DSTs on MPE/iX 6.5 brings many advantages to

HP 3000 customers and paves the way for future capacity expansions in many more areas of MPE/iX. The increase to 65535 DSTs is expected to meet the requirements for years to come and carry MPE/iX forward through its future enhancements.

User-Defined Job Queues and ALTJOB HIPRI Enhancements

*by Jayaram and S Shashikala
Commercial Systems Division*

Enhancements Summary

Previously MPE/iX had one job queue into which all the submitted jobs went before getting launched by the dispatcher. Often this proved to be a hindrance, as a few long jobs could prevent a lot of short jobs from running.

MPE/iX can now have up to 255 user-defined job queues, each having a separate limit on the number of runnable jobs. Users can specify the queue to which their jobs should log on by using the newly added ;jobq= parameter of !job command or !stream command. There will be one default job queue named HPSYSJQ, which will be created during the system bootup.

Three new commands have been added: NEWJOBQ, PURGEJOBQ and LISTJOBQ

- LIMIT, ALTJOB, JOB and STREAM commands now have a new parameter, ;jobq= .
- LIMIT command can now be used to increment or decrement the limit values.
- ALTJOB command has a new option, ;HIPRI to change the priority of the job as HIPRI. This option can be used to override the job limit. SM or OP capability is required for this option.

NOTE The following commands have been changed since MPE/iX 6.0

- NEWJOBQ and PURGEJOBQ commands accept only alphanumeric characters for queue name. A queue name longer than 8 characters will not be accepted.
- STREAM, JOB, LIMIT, ALTJOB commands accept only alphanumeric characters for ;JOBQ= parameter. A queue name longer than 8 characters will not be accepted.
- SHOWJOB #Jnn;JOBQ is a valid command.

New Commands

NEWJOBQ

Syntax:

```
NEWJOBQ qname [:limit=n]
```

The `NEWJOBQ` command creates a new job queue. SM/OP capability is required to execute this command.

`Limit` is the only queue controlling property. The jobs in the queue are sorted by their `INPRI`. In case of a tie for `INPRI`, jobs are sorted by their `INTRO` time.

The global limit takes precedence over individual queue limits. That is, even if a jobq has a slot available, if the overall limit has been reached, jobs have to wait until one of the jobs finishes or the global limit is increased. When a global slot becomes available, the next job is picked from among the eligible jobqs (those which haven't yet reached their individual limits) using the following algorithm:

- Across all job queues, the highest `INPRI` job is selected. In case of a tie for `INPRI`, the one which was introduced earliest is selected. There cannot be a tie in `INTRO` time.
- The job queues do not persist across reboots, unless a `START RECOVERY` is done. Any other system starts will cause the job queues to be deleted and they will have to be created again.
- This command is available in a session, job, or in `BREAK`. Pressing **[Break]** has no effect on this command. This command is not allowed in `SYSSTART`.

Parameters:

qname Name of the queue to be created. Queuename can contain only alphanumeric characters. Queuename can be up to 8 characters in length; longer names will not be accepted. If a queue of this name already exists, an error is returned.

limit Maximum number of jobs that can be allowed in this queue. The `limit` value can be changed using the `:limit [+]n;` `jobq=` command. If omitted, a value of zero is assumed.

Examples:

```
:NEWJOBQ MYJOBQ;limit=100
```

PURGEJOBQ

Syntax:

```
PURGEJOBQ  qname
```

The PURGEJOBQ command deletes a job queue. The queue will be deleted only if it is empty, that is, if no jobs are waiting or executing in the queue. The default system job queue can not be purged. The user must have SM/OP capability to execute the command. This command is available in a session, job, or in BREAK. Pressing [Break] aborts the execution of this command. This command is not allowed in SYSSTART.

Parameters:

qname Name of the queue to be deleted.

Examples:

```
:PURGEJOBQ  myjobq
```

LISTJOBQ

Syntax:

```
LISTJOBQ
```

The LISTJOBQ command allows the user to list all the existing job queues in the system. It displays the queue name, limit, number of jobs in the queue that are in the EXEC state and the total number jobs in the queue, (No. of jobs in the EXEC state + No. of jobs in the WAIT state). This command is not allowed in SYSSTART.

Examples:

```
:listjobq
```

JOBQ	LIMIT	EXEC	TOTAL
HPSYSJQ	3500	1	1
MYJOBQ	100	1	1
MJQ	10	1	2

Modified Commands

LIMIT

Syntax:

```

                                {numberjobs          }
LIMIT [{ + }] [{,numbersessions      }]
      { - }   {numberjobs,numbersessions}
      [ ;jobq= <qname> ]
  
```

The `LIMIT` command allows users with OP capability to change the job and session limits. As a result of this enhancement, it can be used to change the limit value of individual job queues. The `LIMIT` command now accepts a new parameter `;jobq=`. It also accepts '+' or '-' before the numeric values to indicate increment or decrement.

Parameters:

qname Name of the job queue whose limit is to be changed. If the queue name is not specified, `LIMIT` command will alter the global limit values.

+/- If +/- is specified before the numeric value, the corresponding limit value will be incremented/decremented.

Examples:

```

: LIMIT 60 ;jobq=myjobq { Set the limit of MYJOBQ to 60}
: LIMIT -1 ;jobq=myjobq { Decrement the limit by 1}
: LIMIT +1 { increment the global job limit by 1}
  
```

JOB**Syntax:**

```

JOB [jobname[,username[/userpass].acctname[/acctpass]
    [,groupname[/grouppass]]
[;TIME=cpusecs]
    {BS}
    {CS}
[;PRI=      ]

    {DS}
    {ES}

[;INPRI=inputpriority]]
[;HIPRI      ]

[;RESTART]

[;OUTCLASS=[device][,outputpriority[,numcopies]]]

[;TERM={termtype}]
[;PRIVATE]
[;SPSAVE]

[;JOBQ= <qname>]

```

The **JOB** command now accepts a new parameter *;jobq=*. Users can specify the job queue name into which a particular job should logon. If a job queue is specified in the **STREAM** command, then it overrides the name specified in the **JOB** command. If no queue name is specified, the default system job queue is used (**HPSYSJQ**).

Parameters:

qname Name of the queue into which the job should logon.

Examples:

```
:JOB foo,manager.sys;jobq=myjobq
```

ALTJOB

Syntax:

```
ALTJOB [JOB=] {#Jnnn}  
        {#Snnn}  
  
[[ ;INPRI=]inputpriority]  
  
[[ ;OUTDEV=]{ldev    }]  
        {devclass}  
  
[ ;JOBQ=qname]  
  
[ ;HIPRI]
```

ALTJOB alters the attributes of waiting or scheduling jobs. ALTJOB now accepts a new parameter *;jobq=*. Jobs waiting or executing in a queue can be moved to other queues. Only users with SM/OP capability can move jobs.

When an executing job is moved, the limit of the target queue is ignored and the job continues to execute in the new queue. A waiting job continues to wait in the new queue if the queue has already reached its limit.

The ALTJOB command now accepts a new optional parameter *HIPRI*. The *HIPRI* option can be used for overriding the system jobfence or the joblimit. User must have System Manager(SM) or Operator(OP) capability to use this option. Also, the *INPRI* and *HIPRI* options may not be specified together.

Parameters:

qname Name of the job queue into which the job is to be moved.
HIPRI *HIPRI* option can be used to change the priority of a job.

Examples:

```
:ALTJOB #jnnn;jobq=mjq  
:ALTJOB #jnnn;hipri
```

STREAM

Syntax:

```
STREAM [filename] [,char]
[;AT = timespec]
[;DAY = {day-of-week      }
        {day-of-month     }
        {day-until-month}  ]
[;DATE = datespec]
[;IN = [days[, [hours] [,minutes]]]]
[;JOBQ = qname]
```

The STREAM command now accepts a new parameter `;jobq=`. User can specify the queue name into which a particular job should go. The name specified overrides the queue name specified in the JOB command.

Parameters:

qname The name of the queue into which the job must logon. If no queue name is specified the default system job queue will be used (HPSYSJQ).

Examples:

```
STREAM FOO; jobq=myjobq
```


SHOWJOB

Syntax:

```

[[#]Snnn      ]
[[#]Jnnn      ]
SHOWJOB [STATUS      ][;JOBQ][;*listfile]
        [SCHED        ]
        [item[;item[;...]]]

```

The SHOWJOB command displays status information about jobs/sessions. It now accepts a new parameter ;JOBQ which indicates the queue name to which the job belongs. If the ;JOBQ option is not specified, the SHOWJOB output is the same as before.

A new field JOBQ is added to the showjob output format. The JLIST field in the showjob output format displays only 5 characters. In addition, the SCHEDULED-INTRO heading changes to SCHED-INTRO, in the showjob output format when the JOBQ option is used.

Examples:

1. :SHOWJOB;jobq

```

                                New field
JOBNUM  STATE IPRI JLIST JOBQ      INTRODUCED  JOB NAME
#J3     EXEC      LP   HPSYSJQ  WED 11:46A  FTPMON,FTP.SYS
#J7     EXEC      LP   SYSMGRQ  WED  5:47P  EMG,MGR.SYSMGR
#S81    EXEC      34           THU 12:17P  MGR.GOPI

```

2. To display a particular job and job queue name, enter

:SHOWJOB #J7;JOBQ

```

JOBNUM  STATE IPRI JIN  JLIST JOBQ      INTRODUCED  JOB NAME
#J7     EXEC      2  LP   SYSMGRQ  WED  5:47P  EMG,MGR.SYSMGR
JOBFENCE= 7; JLIMIT= 10; SLIMIT= 60

```

3. To display all scheduled jobs with job queue name, enter

:SHOWJOB ;JOBQ

```

JOBNUM  STATE IPRI JIN  JLIST JOBQ      INTRODUCED  JOB NAME
#J2     EXEC      10S LP   HPSYSJQ  TUE  2:50P  JINETD,MANAGER.SYS
#S2     EXEC      20  20           TUE  2:53P  KISHORE,MANAGER.SYS
#S4     EXEC      2   2           TUE  3:25P  SHASHI,MANAGER.SYS

```

3 JOBS (DISPLAYED):

```

0 INTRO
0 WAIT; INCL 0 DEFERRED
3 EXEC; INCL 2 SESSIONS
0 SUSP

```

JOBFENCE= 7; JLIMIT= 60; SLIMIT= 60

CURRENT: 5/11/99 15:25

```

JOBNUM  STATE IPRI JIN  JLIST JOBQ      SCHED-INTRO  JOB NAME
#J3     SCHED  8  10S LP   MJQ      5/11/99 19:00 JAY,MANAGER.SYS
1 SCHEDULED JOB(S)

```

NEWCI Command (CI Elimination)

by Scott McClellan
Commercial Systems Division

MPE/iX Release 6.5 supports a new command called `NEWCI`. The `NEWCI` command is similar to the pre-existing `RUN` command in that it is used to execute a program. Every time a program executes, a new process is created. A process can be defined as an “independent instance or execution of a program”. Both the `NEWCI` and `RUN` commands create a new process. The `RUN` command creates the new process as a child of the process that invoked the `RUN` command. In the case of the `NEWCI` command, the new process replaces the MPE/iX Command Interpreter (CI) process for that session. Hence the name of the command. `NEWCI` is used to create a new CI process that replaces the old CI process.

The sole purpose of the `NEWCI` command is to reduce the number of processes by the system per session logon. It is intended for use on heavily loaded, high-end systems. The target customer environment is discussed in more detail below. Other than saving one process per user, the `NEWCI` command does not provide any additional functionality over the `RUN` command.

Background Information

The maximum number of concurrent processes is fixed for a given release of the operating system. Hewlett-Packard periodically makes enhancements to the MPE/iX to support more concurrent processes. Nevertheless, on high-end heavily loaded systems, processes can be a relatively scarce resource - potentially limiting the maximum number of users on that system. The maximum number of concurrent processes has not been increased in Release 6.5 and it remains 8190.

Processes are used anytime a program is executing. The operating system itself is a collection of programs (sometimes referred to as system programs). Every time a user logs on to the system several processes are used, a minimum of two processes for system programs, and additional processes for any programs executed by the user. One of the processes that is associated with every user's logon is the Command Interpreter (CI) process. The MPE/iX operating system comes with a standard program which reads and executes commands for the user. The standard MPE/iX CI is found in the program file `CI.PUB.SYS`.

In production environments, most system administrators set up their systems so that users have no direct access to the standard MPE/iX CI. This is typically done by creating some kind of logon UDC that will automatically execute the appropriate application for the user. Such logon UDCs typically disable the `BREAK` key (using `OPTION NOBREAK`). The end result is a production environment that is easier for users (because they do not have to be trained on how to find and run programs, set up file equations, etc), and easier for system administrators to keep secure (because users have no access to CI commands).

Logon UDCs typically use the `RUN` command to execute the user's application program. The `NEWCI` command provides an alternative to the `RUN` command which will result in saving one process per user. The critical difference between the `RUN` command and the `NEWCI` command is that the `RUN` command creates the specified program as a child process which

executes underneath the caller (which is generally the MPE/iX CI). By contrast the NEWCI command will cause the specified process to replace the MPE/iX CI process. The MPE/iX CI is terminated - thus one process is saved.

Syntax

```
NEWCI progfile[, [ " ]entrypoint[ " ]]
    [ ;NOPRIV ]
    [ ;LMAP ]
    [ ;DEBUG ]
    [ ;MAXDATA=maxstack ]
    [ ;PARAM=parameternum ]
    [ ;STACK=stacksize ]
    [ ;DL=dlsize ]
    [ ;NMSTACK=nmstacksize ]
    [ ;NMHEAP=nmheapsize ]

    { *formaldesig }
    [ ;STDIN=[ { fileref } ] ]
    { $NULL }

    { *formaldesig }
    [ ;STDLIST=[ { fileref[ ,NEW ] } ] ]
    { $NULL }

    { BS }
    [ ;PRI={ CS } { # } ]
    { DS }
    { ES }

    [ ;LIB={ P } ]
    { S }

    [ ;XL="library[ , ... ]" ]
    [ ;NOCB ]
    [ ;INFO=quotedstring ]
    [ ;UNSAT=[ " ]unsatproc[ " ] ]
```

Parameters

All of the parameters for the NEWCI command are identical (syntactically and semantically) to the RUN command. For this reason, we did not repeat the full description of all the parameters in the on-line help text for NEWCI. Users are referred to the help text for RUN for complete descriptions of all the parameters except the *progfile* parameter. The *progfile* parameter is described in detail in the NEWCI help text (as shown below).

Description of **PROGFILE** parameter from help text for **NEWCI**

The only required parameter is *progfile*. If any additional parameters are specified, they will override the default parameters that the creator of the program established, but only for that particular execution of the program.

progfile The name of the program file to be executed. The file name can be specified in either MPE/iX syntax (the default) or HFS syntax (where the name must begin with either a . or a /). If the file name is specified using MPE/iX Syntax, it does not have to be fully qualified. File names which are not fully qualified will be qualified based on the current logon environment. The file name may be redirected with a file equation. If the file name is expressed using HFS syntax (e.g.: the name begins with a . or a /), and the file name refers to a file outside of MPE/iX name space (e.g. the file is not in an MPE/iX group/account), then some restrictions apply.

The following restrictions are placed on programs outside of MPE/iX name space:

- The program cannot be linked with PM, MR, or DS capability. Programs linked with these capabilities will not load.

- If the program is linked with PH capability, then users must have PH capability to load the program.
- CM Programs cannot be loaded from the HFS directory.

For detailed information on any of the other NEWCI parameters, refer to the RUN command help text.

Operation and Use of New Command

The syntax for the NEWCI command (and all of the parameters) is identical to the RUN command. The behavior of the NEWCI command differs from the RUN command in several important ways:

- NEWCI replaces the calling process with the specified one. The calling process is terminated.
- NEWCI is executable only from a session.
- NEWCI is only executable from the root CI process.
- NEWCI will cause **BREAK** to be disabled (see OPERATION below).
- NEWCI is intended for a very specific use (see OPERATION below), RUN is a more “general purpose” command.

The NEWCI command can only be executed from a session; it is not allowed in a job.

The NEWCI command is only executable from the root CI process, which is sometimes referred to as the usermain process. Normally, NEWCI is executed from the root CI process when executing a logon UDC. If a NEWCI command is executed, then the newly created process replaces the existing root CI process and becomes the new root CI process for that session. The NEWCI command can be executed programmatically provided the caller is executing in a session and the calling process is the root process for the session.

The NEWCI command is not executable in break mode. If a user runs a program (via the RUN command) and then hits **BREAK**, then that user’s session is in break mode. The SHOWME command displays a message indicating a session is in break mode. When in break mode, if the user tries to execute another program via the NEWCI command (or the RUN command), then the user will be asked if they wish to abort the program that is current running. If the user answers “YES” then the current program will be aborted and the new program will execute. In the case of NEWCI, the new program will replace the existing CI process.

The NEWCI command disables break handling for the session. Break remains disabled even if the program calls FCONTROL to re-enable break. The only way to re-enable break for the session is to execute “newci ci.pub.sys”. This causes the MPE/iX Command Interpreter to replace the current CI, and CI.PUB.SYS re-enables break. CI.PUB.SYS also re-executes logon UDCs which may in turn disable break via the option nobreak feature.

NEWCI effects various session resources as follows:

- Any file equations which were set prior to :NEWCI are preserved.
- Any variables which have been set prior to :NEWCI are preserved.
- Any temporary files which have been created prior to :NEWCI are preserved.

- Any UDC files which were cataloged prior to the NEWCI are still cataloged (and thus are executable via the HPCICOMMAND intrinsic).
- DSLINES which were opened prior to :NEWCI will be closed.

NEWCI CI.PUB.SYS

The NEWCI command can be used to replace the current CI process (which may not be the standard MPE/iX CI if a NEWCI command was previously done) with the standard MPE/iX CI. This is done by executing the command “newci ci.pub.sys”. If CI.PUB.SYS is executed using the NEWCI command, the behavior is slightly different than if it is executed using the RUN command. The behavior of the MPE/iX CI is identical to its behavior at logon time.

EXAMPLES:

NEWCI from a Logon UDC

The most common usage of NEWCI is from a logon UDC (most commonly an OPTION NOBREAK logon UDC). Many system managers set up their users such that a logon UDC automatically executes the appropriate application program at logon time. Consider the following example logon UDC:

```
dologon
option logon nobreak
file progdata=datafile.pub.payroll
newci payroll.pub.payroll
*****
```

NEWCI Programmatically

The NEWCI command is programmatically executable, but only from the “root” CI process. Some programs allow users to interactively enter CI commands (by convention CI commands are generally prefixed with a : character). Most HP products/utilities allow CI commands to be entered interactively. In the example below, the NEWCI command is executed programmatically from within TDP. In this example, TDP is the root CI (it became the root CI when the first NEWCI command was issued).

```
:hello mgr.paryoll
:newci tdp.pub.sys
TDP/V (A.05.05) HP36578 Editor (c) COPYRIGHT Hewlett-Packard Co. 1993
/:newci payroll.pub.payroll
CORPORATE PAYROLL (Version A.00.00)
CMD>
```

NOTE In the above example the “:newci payroll.pub.payroll” command would not be allowed if the RUN command had been used in place of the NEWCI command to load TDP.

NEWCI CI.PUB.SYS

The NEWCI command can be used to replace the current CI process with the standard Hewlett-Packard Command Interpreter (CI.PUB.SYS). When CI.PUB.SYS is executed using the NEWCI command it functions exactly the same as it does at logon time (see table above).

```
:hello mgr.payroll
HP3000 Release: X.60.16 User Version: X.60.16 WED, JUL 14, 1999, 9:21 PM
MPE/iX HP31900 C.16.01 Copyright Hewlett-Packard 1987. All rights reserved.
** System Welcome Message...
:newci payroll.pub.payroll
CORPORATE PAYROLL (Version A.00.00)
CMD> :newci ci.pub.sys
** System Welcome Message...
:
```

NOTE This example assumes that the application program payroll.pub.payroll allows the user to enter CI commands interactively (prefixed with a :). This is a common feature in MPE/iX application programs - but it does vary from program to program. In this example the command “newci ci.pub.sys” could be executed by either the COMMAND or HPCICOMMAND intrinsic.

Compatibility Considerations for COBOL and C

by Walter Murray
Support Technology Lab

This article discusses compatibility concerns relating to HP COBOL II/iX and HP C/iX on MPE/iX 6.5. These concerns are the result of enhancements to provide better language support for 64-bit integers.

The project to improve support for 64-bit integers involved changes to four products.

- The C/iX compiler was enhanced to provide full support for “long long” integers. Details can be found elsewhere in this Communicator.
- Seven new routines were added to the Millicode libraries, `MILLI.LIB.SYS` and `/lib/libmilli.a`, to perform highly optimized arithmetic operations on 64-bit integers. The entry points for these new routines are `$$mul2U`, `$$mulo2U`, `$$mul12U`, `$$div2I`, `$$div2U`, `$$rem2I`, and `$$rem2U`.
- The C/iX Library was enhanced with two new functions to convert from strings to 64-bit integers. These are `strtoll()` and `strtoull()`. In addition, the `printf()` and `scanf()` families of functions were enhanced to recognize an optional `ll` (ell ell) modifier to specify that the corresponding argument is a long long integer.
- The COBOL II/iX compiler was updated to generate code that takes advantage of the new 64-bit Millicode routines.

Forward Compatibility Not a Problem

Hewlett-Packard strives to maintain forward compatibility from one release of MPE to the next. MPE/iX 6.5 is no exception. Applications developed on previous releases of MPE/iX should continue to run without needing to be recompiled or relinked. Furthermore, whenever an application is recompiled on MPE/iX 6.5, it should not be necessary to make any changes to the source code or to the procedures for compiling, linking, and running.

Backward Compatibility Not Guaranteed

In general, HP does not guarantee that an application can be developed on one release of MPE/iX and then executed on an earlier release without change. Although code that is compiled using the latest release can often be moved successfully to a system on an older release, doing so is generally not supported.

As is explained below, moving executable programs and executable libraries from MPE/iX 6.5 to a pre-6.5 system probably works in many cases. However, there are situations where special procedures might be required if you want to avoid having to recompile or relink on the earlier system.

It isn't possible to enumerate all of the situations in which problems might be encountered when moving backward to a pre-6.5 system. If you encounter difficulties in doing so, you may be able to find workarounds. Remember, though, that in some situations it may be necessary to recompile in order to move an application from MPE/iX 6.5 to an earlier system.

COBOL/iX

Backward compatibility for COBOL II/iX executable programs (NMPRGs) and executable libraries (NMXMLs) should not be a problem. If your COBOL code generates calls to any of the new Millicode routines, those routines will be copied from the Millicode library and bound into your program or XL. The absence of the latest Millicode library on the target machine is not a problem, because the Millicode library is a relocatable library and is not accessed at run time.

However, if you compile a COBOL II/iX program on an MPE/iX 6.5 system and try to link it on a pre-6.5 system, you may have unresolved externals because of calls to the new Millicode routines. These unresolved externals might be reported at link time, but it is possible that the problem would not show up until load time, with unrelated errors reported by the `RUN` command. The recommended procedure is to link on the 6.5 system, where the latest Millicode library is available.

C/iX

As with COBOL II/iX, the C/iX compiler may take advantage of the new Millicode routines, even in programs that don't use the new "long long" 64-bit integer data type. If you compile and link on a 6.5 system, you will probably not have difficulty moving a C/iX executable program (NMPRG) to a pre-6.5 system. But if you compile on a 6.5 system and try to link on a pre-6.5 system, you may encounter unresolved externals or unrelated load-time errors, especially if you use 64-bit integers.

There are additional considerations to be aware of if you use the 64-bit integer functionality of the C/iX Library. If you do so, and if you want to move your code to a pre-6.5 system, you need to be sure to link with `LIBC.LIB.SYS`, the relocatable library that contains the entire C/iX Library. If you link with `LIBCINIT.LIB.SYS` or `LIBCSHR.LIB.SYS` instead, and then move your program to a pre-6.5 system, there are two problems you are likely to encounter.

First, if you have called `strtoll()` or `strtoull()`, you will have unresolved externals, because these functions do not exist in the C/iX Library in `XL.PUB.SYS` on pre-6.5 systems.

Second, you will experience strange behavior if you use 64-bit integers with the `printf()` and `scanf()` families of functions. The pre-6.5 versions of these functions do not recognize the `ll` modifier for long long integers and will produce incorrect results.

Not Supported, but Might Work

In summary, it is not supported to compile code on any release of MPE/iX and then move it to an earlier release, even though it often works.

If you try to move code from an MPE/iX 6.5 system to a pre-6.5 system, if it doesn't work, the following are the most likely symptoms.

- You get unresolved externals for `$$mul2U`, `$$mulo2U`, `$$mul12U`, `$$div2I`, `$$div2U`, `$$rem2I`, or `$$rem2U`. If the unresolved externals go undetected, you get strange, unrelated errors at load time. This could happen with COBOL II/iX or C/iX.
- You get unresolved externals for `strtoll()` or `strtoull()`. This could happen if you use these new C/iX Library functions for converting 64-bit integers.

- Your `printf()` and `scanf()` calls produce incorrect results. This could happen if your code uses the new functionality for handling 64-bit integers with these functions.

If you do try to move a COBOL or C application from a 6.5 system to a pre-6.5 system and it doesn't work, you probably can work around the problem. As a last resort, you may have to recompile on the pre-6.5 system.

C/iX Supports 64-bit Integers

*by Walter Murray
Support Technology Lab*

In response to a number of customer requests, we have enhanced the C/iX compiler, the C/iX Library, and the MPE/iX Millicode library to provide full support for a 64-bit integer data type. This article provides the information you need in order to declare, operate on, and perform input/output operations on 64-bit integers in C.

All of the changes have been made in such a way as to ensure forward compatibility to MPE/iX 6.5 from pre-6.5 systems. All existing C/iX applications should work the same on 6.5 as they did on earlier systems, whether you recompile and relink on 6.5, or whether you bring your executable programs and libraries forward to 6.5 without recompilation.

However, there are some backward compatibility issues to be aware of if you compile C code on MPE/iX 6.5 and want to move it to a pre-6.5 system. Please see the article "Compatibility Considerations for COBOL and C", elsewhere in this Communicator.

The 64-bit integral type is referred to in C/iX as "long long". The long long data type is an extension to ANSI C. If you compile with the CCXL command and want to use 64-bit integers in ANSI mode, you must enable ANSI extensions by using the compiler option "+e" or "-Ae". If you compile with the c89 command, enable ANSI extensions by specifying "-Wc,+e".

To declare a signed 64-bit integer, use the type specifier "signed long long int", where "signed" and "int" are optional. To declare an unsigned 64-bit integer, use the type specifier "unsigned long long int", where "int" is optional. Note that, for compatibility reasons, the type "long int" remains 32 bits, the same as type "int".

A 64-bit integer is stored in 8 bytes, 8-byte aligned, with higher-order bits stored in lower-addressed bytes. (The high-order 32 bits are stored in the "left" word.)

A signed long long int can represent integers in the range -9223372036854775808 to 9223372036854775807. An unsigned long long int can represent integers in the range 0 to 18446744073709551615.

To write a signed long long integer constant, whether decimal, octal, or hexadecimal, you must use a suffix of "LL". Likewise, an unsigned long long integer constant must be written with a suffix of "ULL". (Suffixes may be written in upper, lower, or mixed case.)

Within your C code you may use variables of type long long the same as you would use variables of any other integral type. For example, a variable of type long long can be used as an operand in an arithmetic or relational expression, as a function argument, parameter, or return, as a subscript, etc. You can declare arrays, structure members, and bit-fields of type long long. You can take the address of a long long, use a cast between long long and other types, and initialize a long long at the point where it is declared.

For converting to and from 64-bit integers, the scanf() and printf() families of functions have been enhanced to permit "ll" (ell ell) as a modifier to a conversion specifier of "d", "i", "o", "u", "x", or "X". There are also two new functions, strtoll() and strtoull(), which are the 64-bit equivalents of strtol() and strtoul().

The following sample program is an example of how to declare, use, and print a 64-bit integer using C/iX.

```
/* Compile on MPE/iX 6.5.          */
/* With CCXL use "-Aa +e" or "-Ae". */
/* With c89 use "-Wc,+e".        */
#include <stdio.h>
main(void)
{
    long long bignumber = 1234567654321LL;
    bignumber = bignumber * 1000000 + 123;
    printf("bignumber = %lld\n", bignumber);
}
```

FTP/iX Support for Large Files

by R. Prakash
Commercial Systems Division

The enhancement “FTP/iX support for Large Files” which is part of MPE/iX Release 6.5 provides the capability for FTP on MPE/iX to transfer files of size greater than 4GB. This enhancement is an offshoot of the support for Large Files in MPE/iX 6.5.

Until MPE/iX 6.0, you could transfer files up to 4GB in size through FTP. With this new enhancement, FTP/iX is now capable of transferring files beyond 4GB also.

Features

The MPE/iX Large File types which can be transferred through FTP/iX are:

- Fixed ascii flat Large Files
- Fixed binary flat Large Files
- KSAM64 Large File (newly defined KSAM file type)

It is also possible to append to a fixed ascii flat Large File.

It can transfer Large Files only between compatible peer MPE/iX systems which are enhanced to support Large Files.

User Interface

The syntax for FTP/iX commands to transfer Large File remains the same as for non-Large Files. There are no changes in the command syntax given by the user. There are few minor changes with respect to user interface output to indicate to the user that the FTP/iX server is large-file capable.

When FTP/iX tries to talk to a Large File compatible server, the following is the snap shot:

```
-----<Start>-----  
Connected to hanuman (15.70.188.86). (FTPINFO 40)  
Name(manager): manager.sys  
---> USER manager.sys  
230 User logged on  
---> SYST  
215 MPE/iX LF system type. <----- Note 'LF' here  
Remote system type is MPE/iX  
---> SITE MPE/iX FTP Client [A0010001]  
200 MPE/iX command ok.  
---> TYPE I 200  
Type set to I.  
  
-----<End>-----
```

It prints “MPE/iX LF system type” instead “MPE/iX system type”.

FTP/iX cannot transfer a Large File if it is pre-6.5 Client/Server. The following is one of the snap shots when it tries to transfer a Large File to a non-compatible server.

```
ftp> put ksm64001  
Server Not compatible to transfer largefile  
Data Transfer Request Failed. (FTPERR 13)  
ftp>
```

Installation

There are no changes in the installation procedures.

4 Data Center Management

This chapter contains technical articles that provide information about new or enhanced components of data center management in the MPE/iX 6.5 release. The following articles are included in this chapter:

- Enterprise Management Solution
- Support Tools Manager (STM)
- Enhanced Message Course Template for HP 3000 Agents in ITO
- HP Secure Web Console
- TurboSTORE Support for DLT
- DLT4000/DLT7000 Differential Tape Support on MPE/iX
- DLT8000 Differential Tape Support on MPE/iX
- TERMDSM: User Interface Change Under System Diagnostics
- HP SureStore E Disk Array 12H on the HP 3000
- HP SureStore E Disk Array XP256

Enterprise Management Solution

*by Saurabh Kelkar
Commercial Systems Division*

The Current Need

As data centers become larger, the complexity of managing them increases manifold. Today's challenge is to manage distributed, heterogeneous IT environment that comprises various systems. The level of sophisticated management has to be extended towards business critical applications and service management, where scalability, reliability, robustness, and optimized down time are the key factors. At the same time, we have to keep the Data Center Management simple enough such that the customer spends a minimal amount of effort on developing the skill set to use them.

Today there is a whole range of Enterprise Management solutions available in the market. When we look at the entire collection of the solutions, we realize that we have some opportunities to fill some gaps in terms of offerings that are available on the HP 3000 platform. The HP 3000 needs to get integrated seamlessly in today's EM setup to enable the users to manage it with little or no extra overheads.

Keeping this need in mind, we formed an Enterprise Management Solutions' group, who can focus on fulfilling the user communities' needs in this area. We would like to let you know the roadmap for the same, that is, the direction we intend to move.

Recent Offering

Among our recent offerings, we are pleased to provide the Legato Backup Solution and free availability of EasyTime on the Fundamental Operating System (FOS).

The Legato Backup solution enables the HP 3000 to be backed up in a heterogeneous machine environment. Today we have many customers whose data centers are spread across great distances. Not only that, customers have a plethora of machines from different vendors for different tasks in their data centers. Legato Backup solution becomes an ideal solution for such an environment.

EasyTime allows users to work on the HP 3000 effectively without needing to know much about the Command Interpreter of the MPE/iX OS. It allows the data center manager to make group of commands that are easy to use and tailored to his need.

Added to this, we have made an effort to make the MPE/iX programming documentation better and easily available – for example, we have made some of the documentation available on web for easy accessibility.

Short Term Plans

In the next 1-2 years, we plan to strengthen our basic set of offerings. A very near future inclusion is providing support for the HP Secure Web Console (SWC). The inclusion of HP SWC in the list of devices supported by the HP 3000 means that the users have 'Location Independence' for accessing the HP 3000 system console. This is, of course, coupled with a high level of security to provide a secure access to the system's console. Please see the separate article on the SWC for more information.

The data center manager needs to manage all his machines efficiently. It has become a challenge to manage a heterogeneous geographically dispersed data center cost effectively and efficiently. We already have the IT/Operations Agent available on MPE/iX. We have enriched the default environment for the agent very recently. Now the ITO's default settings are tuned better with the MPE/iX environment. We believe this should result in operational convenience for the users. Please see the separate article on the New Default Templates on the IT/O on MPE/iX for more information.

We are planning to make the Omniback Agent available on MPE/iX shortly. With the availability of the Omniback 3.1 Agent, we essentially provide the HP 3000 user community a powerful set of back-up solutions to choose from. We think this adds a great degree of flexibility for the users to select their choice of backup solution.

To integrate the HP 3000 in a data center for performance measurement purposes, we are considering providing a powerful performance solution like the Measureware Agent available on the HP 3000. We may first provide an interim work-around to allow the existing Scope Collector to smoothly send the measurement data to the central PerfView console.

Long Term Plans

In the long term, we plan to extend the strength with new and additional capabilities.

With the constantly increasing demand to process more and more data, users are going to need a solution that allows increased availability of data, increased backup performance, simplified management of growth, etc. We think our users would eventually move to a SAN environment. As mentioned earlier, in the long term, we intend to enable the HP 3000 to integrate seamlessly in the SAN environment.

We plan to start putting various building blocks for SAN support into the base platform. The list of such 'building blocks' includes solutions like - system level integration of key user applications, simplified/automatic device configuration, simplified database maintenance, web-based centralized system management, improved (and faster) back-up for backing up terra bytes of data efficiently, and even tools for simplified and quick resource planning.

Summary

We know that the data center management needs are going to grow. The need to handle large amounts of data, the need to have highly available data, the need for increased performance of backup, the need to have a better fault tolerant storage system, the need to have a scalable storage system etc. are some of the future needs.

We plan to continue making the data center management an easy and complete job for HP 3000 customers such that the skill set needed to manage the HP 3000 in a data center is going to be very small. Our goal is to create a Virtually Operator-less environment.

As we get more clarity on the market and our customers' needs, we will proactively change/tweak our plans to better reflect the needs of the customers. With this strategy, we will ensure that we are putting our resources on the right solutions, i.e. the solutions that our customers want most.

Support Tools Manager (STM)

by Bart Anderson
Diagnostics and Support Tools

With MPE/iX 6.5, the Support Tools Manager (STM) replaces the `sysdiag` diagnostic system on HP 3000 computer systems. STM has been the diagnostic platform for HP-UX since 1996.

STM was ported to MPE/iX to provide a common diagnostics platform for MPE/iX and HP-UX systems. This common platform reduces the cost to support new systems and peripherals, and makes it possible to bring new tools to MPE/iX 6.5 and beyond. STM is not supported on any previous MPE/iX release.

NOTE For STM to run, networking to be configured (`NMMGR`) and enabled. (STM requires sockets.) Users will not be able to interact with the Support Tools Manager until networking has been started.

The STM package provides a set of support tools for troubleshooting and monitoring computer system hardware, including CPUs, memory, interface cards, mass storage devices, etc.

Upon startup, STM can provide a system map of all of the hardware devices on your system. Using this system map, you can select devices, receive information, or execute on-line support tools.

STM is distributed on the Fundamental Operating System (FOS) tape on MPE/iX 6.5 and subsequent releases. When previous releases of MPE/iX are updated to 6.5, the `sysdiag` diagnostic system is removed and STM is installed.

In the first phase of the release, STM for MPE/iX includes:

- The diagnostic platform.
- Information tools, expert tools, and firmware download tools for SCSI peripherals.
- Information tools for CPUs and memory.
- Expert tool for memory.
- Existing `sysdiag` diagnostics that are “wrapped” for STM.
- Logtool for I/O and hardware-related logs (based on the HP-UX version of logtool).
- A second, standalone version of logtool similar to the existing `sysdiag` logtool. This version would handle system logs (e.g. auditing, security, etc.) not covered by the STM version of logtool.

In later phases, STM on MPE/iX is planned to include verifiers, exercisers and other tools.

If you are already familiar with STM on HP-UX, you will feel at home with STM on MPE/iX. However, there are some differences to be aware of:

- Different distribution and installation process.

- Different file structure.
- No xstm (GUI) interface for MPE/iX.
- MPE/iX version initially contains a subset of the tools on the HP-UX version.
- An additional standalone version of logtool to handle MPE/iX system logs.
- The first release of STM on MPE/iX is based on the A.12.00 version of STM (June 1998, IPR 9806).

To learn how to use STM, refer to “Support Tools Manager (STM) for MPE/iX” in the *MPE/iX 6.5 System Software Maintenance Manual*.

To learn more about STM, see the “Systems Hardware, Diagnostics and Monitoring” section of Hewlett-Packard’s online documentation Web site at:

<http://www.docs.hp.com/hpux/systems/index.html>

This site provides manuals, tutorials, FAQs, and reference material on STM and other support tools.

Enhanced Message Source Template for HP 3000 Agents in ITO

*by Shirin Mendonce
Commercial Systems Division*

Introduction

In a new ITO setup a default set of message source templates is installed on the Management Server for every managed node platform supported by ITO. These default templates are downloaded from the Management Server at the time of Agent installation onto the managed nodes. The purpose of these default templates is two fold:

- To expose the power of templates and serve as examples for further template creation and fine tuning.
- To serve as a ramp during the initial post installation period for maximum Return on Investments(ROI).

The default template implementations for the different managed node platforms do not all have the same level of functionality. Some platforms have advanced sets of default templates while some have very little capability. The HP 3000 platform (whose OS is known as MPE/iX) had a very basic default template set. With a view to better help management of environments consisting of HP 3000 systems, a new enhanced default template set has been created.

The default message source templates for the MPE/iX platform can be categorized into the following areas:

- OS Monitoring
- Network Monitoring
- Database Monitoring
- Console Monitoring

OS Monitoring

The OS monitoring templates can be classified as:

- Disk Drive Monitor
- Volume Set Monitor
- Job and Process Monitor
- Monitor for other monitors
- Printer Outfence Monitor
- Print Queue Length Monitor
- Pending Console Reply Monitor
- Message Queue File Monitor

Disk Drive Monitor

This template monitors the current status of the disk drives on MPE/iX systems. The executed script, `DEVINFO`, retrieves the disk drive status information and writes to the file `DEVDATA`. Messages to the Management Server are generated for the following statuses of the disk drives:

STATUS	MEANING
UNKNOWN	The system cannot recognize the volume label on the disk
SCRATCH	The volume has been erased through use of disk or mountable volume.
LONER	The volume is a member of a volume set whose master is not mounted, or is a master with the same volume set name as one that is already mounted.
MASTER	The volume is a recognized master volume of a volume set
MEMBER	The volume is a recognized member volume of a volume set
PENDING	The partner disk failed to mount.
DISABLED	The volume is not available to the system due to disk failure

Volume Set Monitor

This template monitors the utilization of volume sets and volume members for MPE/iX systems. The executed script, `VOLINFO.MONITOR.OVOPC`, calculates the current utilization of PERM space, FREE space, and TRANSIENT space in both sectors and percentage used for each volume set on the system and for each volume member within the set. The information is written to the file `VOLDATA.PUB.OVOPC`. `VOLINFO` monitors disk space utilization by volume and volume set. The current utilization is compared against thresholds for volume sets and volume members defined within IT/Operations.

Job and Process Monitor

This monitor verifies if an MPE/iX background job or process is running. The templates have been preconfigured to monitor some of the commonly found background jobs and processes such as `FTPMON`, `JINETD` and `SCOPEJOB`. The monitor program is `PROCMON`.

Monitor for other Monitors

This monitor checks to see if any background job or process monitors are currently disabled. This template is intended to be used as a reminder that some background jobs are not being monitored because they are disabled. This can be quite normal whenever this is a planned shutdown of background processes or jobs such as before a system backup is performed.

Printer Outfence Monitor

This template monitors the current value of the printer outfence. If an LDEV is specified, it returns the value for that printer. If no printer LDEV is given, the global outfence value

is returned. The threshold can be defined as increasing or decreasing depending on how the outfence is normally used for the given device.

An outfence that is higher than expected will result in spoolfiles being deferred that would otherwise print. If the outfence is lower than expected, spoolfiles that should have been deferred are printed. The monitor can be used to monitor either of these conditions.

Printer Queue Length Monitor

This template monitors the number of pending spoolfiles for a given spooled printer. It is useful when a printer is expected to be available to dynamically print a report, such as with a workgroup printer. A large number of pending spoolfiles may signify problems with the printer or that a large print job may be overusing the printer. This can have an impact on service management agreements if responsive printer services are an objective.

Pending Console Reply Monitor

This template monitors the number of outstanding or pending console replies.

While there is a `CONSOLE` template condition for console replies, this monitor provides notification that replies may not have yet been given a response. This template serves as a useful reminder if the initial `REPLY` event is acknowledged without the `REPLY` being satisfied. The severity is higher as more replies are pending.

Message Queue File Monitor

This template is provided as an example for monitoring the number of records contained in an MPE/iX message file. This template is used when the number of records contained in the message queue may result in a delay in process communication, that is, it represents a backlog condition.

Networking Monitoring

These templates monitor the status of the Network Interfaces and the Network Services.

Networking Interface Monitoring

There are three network interfaces which have been found to be the most commonly used ones that have been preconfigured for monitoring:

- LAN interface
- LOOP interface
- FDDI interface

The monitoring is done by a monitor program called `NISTATUS` which verifies that the network interfaces are started and running at intervals of time. The status of the network interfaces is sent to the Management Server.

Network Services Monitoring

This template monitors the status of the NS Network Services. The monitor program `NSSTATUS` returns the status of the Network Services to the Management Server.

Database Monitoring

This template monitors the capacities and database flags for Image/SQL databases. The executed script, `DBINFO.MONITOR.OVOPC`, retrieves the database information and writes it to the file `DBSTATUS.PUB.OVOPC`. The input file, `DBLIST.PUB.OVOPC` needs to be created on each MPE/iX managed node that will be assigned this template. The `DBLIST.PUB.OVOPC` should contain the fully qualified root file names for each Image database to be monitored.

Console Messages

These templates provide for the interception of MPE/iX system console messages. The conditions defined within these template address the following systems management areas:

- Volume Management
- Security
- Backup / Tape Volumes
- OpenView System Manager

Volume Management

These console messages get generated whenever there is a fault in the Volume and Disk Management on MPE/iX are intercepted and received by the Management Server. The messages that have been preconfigured for interception are:

Mirrored volume partner missing messages

Occur when an LDEV's partner does not mount and the LDEV is placed in `PENDING` state as a result.

Disk array disabled mechanism messages

Occur whenever a mechanism in a disk array has suffered a failure.

Mirrored volume disabled messages

Occur when a mirrored disk device has failed after successfully mounting.

Security

These console messages are related to user logons. The messages that come under this category are:

- Invalid logons due to bad password
- Invalid logons due to missing parameters
- User logons with SM capability

- User logons on LDEV 20 which is the console
- User logons on LDEV 21 which is the remote console
- Logons on LDEV 10 which are from jobs

Backup/Tape Volumes

Messages that come under this category are:

- STORE Bad media messages
- STORE Completion status messages
- Tape not ready messages
- Tape mount request pending messages

Open View System Manager

All OpenView System Manager messages generated by virtue of their being instrumented in applications come under this class. This template has been accommodated with a view to help customer smoothly transition from OpenView System Manager to ITO without disturbing their existing setup.

HP Secure Web Console

*by Umakanth Rajsekhar C.
Commercial Systems Division*

Product Overview

This module gives an overview of the HP Secure Web Console. (Product No. J3591A) The HP Secure Web Console (SWC) is now available on the HP 3000 for its user community.

The HP SWC is the first product from Hewlett-Packard's new Internet Product family. The product is an easy-to-use device which offers 'location independence' by guaranteeing secure console access and machine administration from either Netscape or Internet Explorer web browsers on any PC or UNIX workstation client across the network. The HP SWC also offers a flexible and comprehensive migration to a security framework enabling strongly regulated remote access to the company's business servers.

The HP Secure Console's main purpose is to allow `manager.sys` level system network administrators easier access to HP 3000 servers, regardless of their distance from the HP 3000 machine. The main objective is to be able to perform the same level of debug and recovery operations, just like using a local HP terminal, but in this case completely remotely via the LAN and from any WEB browser running on a PC or UNIX client available on the Network.

Installation

The installation of the HP SWC is very straight forward.

The typical installation of the HP Secure Web Console, after completion of the initial configuration and set-up, is to install the device near the desired HP 3000 system. Replace the normal HP terminal(700/96..) by disconnecting the RS-232 Console cable connector from the back of the HP terminal and connecting it directly to the HP SWC's RS-232 port. Then once connected up to the LAN using an available 10base-T cable and powering it on, the HP SWC's internal micro-web server can be accessed across the LAN using either the Netscape or Microsoft Internet Explorer web browser applications running on any PC or UNIX workstation client node.

Once connection to the HP SWC's internal micro HTTP webserver is established, a Java application is downloaded to the connected client node, which from then on acts as the interface for all communications between the client and the attached HP 3000 server.

High level of security is provided with "user authentication" being performed for each user accessing the HP SWC. This results in allowing user login for only those users who have valid user access permission.

To make the HP 3000 system administration task easier, the HP SWC also allows up to four authorized users to access the HP 3000 system console while any one of the users having the write permissions. This is especially useful for system management tasks requiring more than one person to work as a team.

The HP SWC brings ease of use and security to system administration. Future release of

the product will be available on the Internet. Software upgrade is easily achieved using a standard FTP server.

Product Structure

J3591A: HP Secure Web Console (includes localized power supply and quick installation guide).

The HP Secure Web Console comes with embedded firmware, including an HTTP micro web server 100% pure Java application, which is provided in a compact box with a small footprint measuring only 180x105x30mm(the same size as a VHS video cassette tape). The J3591A HP Secure Console product includes a country-specific power adapter module along with the device itself which weighs no more than 280g and has an operating specification of 0-40deg C storage.

Product Features

- Cost Effective and secure webbed system management
- Full HP2392 Terminal Emulation
- Plug and play, intuitive configuration
- Zero software installation time (JAVA Applet download)
- Security framework, password digest user authentication
- Multi-user access for efficient administration
- Software upgrade through network via the WEB
- Compatibility with
 - HP 3000
 - HP9000
 - Windows 95, NT4.0 and UNIX clients
 - Netscape & Microsoft Internet Explorer
- No software changes needed on the HP 3000 to use the HP SWC

New Software Available with Future Releases

- Enhanced Security features
 - SSL client & server authentication
 - Confidentially with using encryption on console data streams
- Integration into HP OpenView Network Management
- Enhanced Java features using Java 1.1
 - GUI interface enhancements
 - Reduced download time
 - Increased security applets

Support from the World Wide Web

HP provides product and support information from the World Wide Web to identify and resolve problems. The site can be accessed at:

<http://www.hp.com/go/webconsole>

For further information you can contact:

Umakanth
umak@india.hp.com

OR

Saurabh
saurabh@india.hp.com

TurboSTORE Support for DLT

by Uma Bondada
 Commercial Systems Division

About the DLT Drives

DLT drives offers the following advantages:

- Faster Backups and more reliable backups
- High capacity (up to 70 GB uncompressed)
- Faster Restore and hence faster disaster recovery

TurboSTORE Products and the DLT

The following Turbostore products support DLT without the fast search option:

Product Number	Product Description
31900B	FOS Store
B5151AA	TurboStore/iX II
B5152AA	TurboSTORE/iX 7x24 True Online Backup

About Fast Search on Tapes

The fast search concept is generally used when “restoring” files selectively from the tape. The MPE/iX environment variable `HPFASTSEARCH` should be set to `TRUE` for “fast search” to happen. Without fast search, files are skipped one at a time until the required file to be restored is reached. During fast search, a number equalled to the intervening tape marks or save set marks, between two consecutive files of the “restoreset”, is calculated and skipped at one shot, to go to the next file to be restored. Fast search greatly improves the “restore” time.

Fast Search on DLT Tape Drives

MPE/iX 5.5 onwards the fast search on DLT is supported with the installation of the following patches:

- MPEJXX0 for labelled DLT tapes. (STORE patch)
- MPEJXT7 for unlabeled DLT tapes. (STORE patch)
- MPEKX37 for support of DLT7000 tape drives. (IO patch)

With respect to the different MPE/iX versions, the patches and their versions that need to

be applied are given in the following table:

MPE/iX Releases			
	5.5	6.0	6.5
Patch required for labelled tapes	JXX0(D)	JXX0(C)	No patch required*
Patch required for unlabeled tapes	JXT7(B)	No patch required	No patch required
Patch required for labelled tapes on DLT7000	JXX0(D) & KX37(D)	JXX0(C) & KX37(I)	No patch required
Patch required for unlabeled tapes on DLT7000	JXT7(B) & KX37(D)	KX37(I)	No patch required
* No patch required means that, the patch has been rolled into the base OS.			

Limitation of JXX0

The MPE/iX command “FCOPY” cannot be used on the tape created using the store with this patch. To allow “fcopy”, a new STORE option “FCOPY” is to be used while “storing”. However this tape cannot support fast search while “restoring”. So FCOPY and fast search are mutually exclusive.

DLT4000/DLT7000 Differential Tape Support on MPE/iX

by Larry Nichoalds, CSY Labs
Jim Hawkins, SSD WTEC

Beginning with Release 5.5 Express 7, MPE/iX supports the following stand-alone Differential DLT Tape drives:

- DLT4000 (C6382A/C6383A) - F/W SCSI interface (not a true F/W device see “Product Overview”)
- DLT7000 (C6374A/C6375A) - F/W SCSI interface

Software Requirements

Customers wishing to use these new DLT devices must be running on MPE/iX 5.5 or LATER version of MPE/iX. On MPE/iX 5.5 (Release C.55.00, MPE/iX 31900C.05.07) a patch provides the critical software device driver enhancements necessary for the support of these DLTs. Again, these devices are NOT supported on MPE/iX 5.0 (Release C.50.00, 31900B.79.06) and there are no plans to support these devices on 5.0 in the future.

Hardware Requirements

On HP 3000 S900 systems F/W SCSI Device Support is typically provided via the 28696A NIO F/W SCSI Device Adapter Card. At the time this article was written, there are two other possible F/W SCSI Device Adapter Cards.

NEITHER of these connection options is actually supported:

Unsupported connection Option #1	The 9x9KS “Core I/O Card” A3453-69210 exports a F/W SCSI interface. However, we do NOT, for performance reasons, support or recommend that a DLT be connected via this interface. The assumption is that this exported bus will also be shared with the system’s internal disks and unpredictable and unacceptable performance may be seen when disks and DLT are accessed simultaneously (as during STORE/RESTORE). (Please read additional information in the Product Overview section for details on device performance.
Unsupported connection Option #2	The HP-HSC F/W DA A2969A is not supported on HP 3000 S900 systems. This card requires the existence of “Device Adapter Software” that has yet to be written for MPE/iX.

Physical connection of these devices is accomplished via “Fast/Wide/Differential SCSI” Cables such as C2911A, C2924A, C2925A etc.

NOTE While DLT Libraries are available for purchase, support for these devices is very limited on HP 3000 S900 systems. There are currently no plans for

supporting the DLT Libraries for “normal” MPE/iX operations such as `BOOT`, `FCOPY` or `STORE`. DLT Libraries are supported only when used in conjunction with a third party software package, Legato’s “NetWorker.” Inquiries about DLT Library support should be made directly to Legato.

Product Overview

Product	Standard, Non-compressed Capacity with DLT IV	Average Maximum Data Transfer Rate to Tape	Average Maximum Data Transfer Rate on SCSI Bus	Maximum of Devices per Bus
DLT4000 (SE-SCSI)	20Gbyte	~2Mbyte/sec	~4Mbyte/sec	1
DLT4000 (FW-SCSI)	20Gbyte	~2Mybte/sec	~4Mybte/sec	1
DLT7000 (FW-SCSI)	35Gbyte	~4Mbyte/sec	~9Mybte/sec	1

The performance difference between the DLT4000 and DLT7000 Differential drives is due to data transfer characteristics. Specifically the DLT7000 is a “differential-wide” device, transferring two bytes at a time during the data phase. The DLT4000 (C3690A), though it is connected via a “differential-wide” cable, is actually a “differential-narrow” device, transferring only one byte at a time. From a native (non-data compression) perspective, the DLT4000 is capable of transferring data to tape at up to 2 Mb/Sec while the DLT7000 is capable of transferring data to tape at up to 4 Mb/Sec. When making use of the compression feature (Digital Lempel-Ziv, DLZ) the amount of data sent to the device can be significantly greater than the device’s ability to write to tape, depending on the compressibility of data. However, there is an upper limit on the amount of data that the device can accept and the SCSI bus can transfer. This puts an upper limit on the number of devices that can be driven at peak I/O rates on a Bus. This upper limit is the “Maximum Number of Devices Per Bus.” Of course, these figures are approximate and are reached only on systems with sufficient available CPU resources AND Disk I/O bandwidth.

Both the DLT4000 and DLT7000 drives accept the new CompacTape IV cartridge. When this cartridge is used, the amount of data you can store on a tape can be up to 20.0 GB native, or an expected 40.0 GB of data using the compression mode for the DLT4000. The DLT7000 will hold 35.0 GB native, or 70.0 GB when employing data compression. The actual amount of data stored in compression mode is dependent on the nature of the data being stored. That is, text files tend to be highly compressible while code files are not. Both devices can also use the CompacTape III cartridge. The amount of data you can store on this tape cartridge can be 10.0 GB native capacity or 20.0 GB compressed.

IMPORTANT Tapes created on a DLT4000 can always be read by the DLT7000; however, DLT4000 drives are not capable of reading tapes created on DLT7000s. Since both drives use the same physical media type, shops with a mixture of DLT types will have to manage media created on DLT7000 separately so as not to attempt to read it on a DLT4000. Typically if a DLT7000 tape is put into a DLT4000, the “Use Cleaning Tape” indicator will be illuminated. If a DLT7000 tape is to be over-written by a DLT4000, just ignore the “Use Cleaning Tape” indicator and proceed overwriting the tape (do not clean the drive). The “Use Cleaning Tape” will go off automatically when the tape is overwritten or removed from the drive.

DLT is designed as a high-volume back-up solution. The tape media can endure approximately 500,000 passes and has a shelf life of 10 years. The DLT4000 has a head life of 10,000 tape motion hours and MTBF of 80,000 based on a 100% duty cycle. The DLT7000 has a head life of 30,000 tape motion hours and MTBF of 200,000 hours with a 100% duty cycle (per Quantum’s literature). Unlike DDS drives, you should only clean DLT devices when the cleaning light actually comes on. If you clean a DLT more often than is indicated, you will reduce the life of the heads.

Configuration

Here is an example of what the configuration should look like on the HP 3000 for the DLT4000 and DLT7000.

DLT4000 Example

```
io> ld 90
  LDEV: 90  DEVNAME:          OUTDEV:  0      MODE:
  ID: DLT4000                RSIZE:  128   DEVTYPE: TAPE
  PATH: 4/4.2.0              MPETYPE:  24   MPESUBTYPE: 8
  CLASS: TAPE                DLTTAPE DDUMP
```

```
io> lp 4/4.2.0
PATH: 4/4.2.0                LDEV: 90
ID: DLT4000                  TYPE: TAPE
PMGR: SCSI_TAPE2_DM         PMGRPRI: 10 LMGR:
LOGICAL_DEVICE_MANAGER      MAXIOS: 0
```

```
io> lp 4/4.2
PATH: 4/4.2                  LDEV:
ID: PSEUDO                   TYPE: DA
PMGR: TRANSPARENT_MGR       PMGRPRI: 6
LMGR:                        MAXIOS: 0
```

```
io> lp 4/4
PATH: 4/4                    LDEV:
ID: HP28696A                 TYPE: DA
PMGR: SCSI_DAM               PMGRPRI: 6
LMGR:                        MAXIOS:
```

DLT7000 Example

```
io> ld 91
  LDEV: 91  DEVNAME:          OUTDEV:  0      MODE:
  ID: DLT7000                RSIZE:  128   DEVTYPE: TAPE
  PATH: 4/4.3.0              MPETYPE:  24   MPESUBTYPE: 8
  CLASS: TAPE                DLTTAPE DDUMP
```

```
io> lp 4/4.3.0
PATH: 4/4.3.0                LDEV: 90
ID: DLT7000                  TYPE: TAPE
PMGR: SCSI_TAPE2_DM         PMGRPRI: 10 LMGR:
LOGICAL_DEVICE_MANAGER      MAXIOS: 0
```

```
io> lp 4/4.3
PATH: 4/4.3                  LDEV:
ID: PSEUDO                   TYPE: DA
PMGR: TRANSPARENT_MGR       PMGRPRI: 6
LMGR:                        MAXIOS: 0
```

```
io> lp 4/3
PATH: 4/3                    LDEV:
ID: HP28696A                 TYPE: DA
PMGR: SCSI_DAM               PMGRPRI: 6
LMGR:                        MAXIOS: 0
```

Announcing DLT8000 Differential Tape Support on MPE/iX

*By Tina Wang, CSY Labs and
Dave Black, MPE/iX Expert Center*

MPE/iX now supports the following Differential DLT Tape drive (this is in addition to the previously announced DLT4000 and DLT7000):

DLT8000 (C6378A) F/W SCSI interface	Stand-Alone
DLT8000 (C6379A) F/W SCSI interface	Rack-Mount

Software Requirements

Customers wishing to use this new DLT device must be running on either MPE/iX 5.5 or 6.0 (and later releases):

- On 5.5 (Release C.55.00, MPE/iX 31900C.05.08, with patch MPEKXU0)
- On 6.0 (Release C.16.01, MPE/iX 31900C.16.01, with patch MPEKXU0)

MPEKXU0 allows the software device driver to recognize the vendor ID or the support of the DLT8000.

The DLT8000 is not supported by Predictive.

Hardware Requirements

On HP 3000 S900 systems F/W SCSI Device Support is typically provided via the 28696A NIO F/W SCSI Device Adapter Card. At the time this article was written there are two other possible F/W SCSI Device Adapter Cards. NEITHER of these connection options is actually supported:

Unsupported connection Option #1

The 9x9KS Core I/O Card A3453-69210 exports a F/W SCSI interface. However we do NOT, for performance reasons, support or recommend that a DLT be connected via this interface. The assumption is that this exported bus will also be shared with the system's internal Disks and unpredictable and unacceptable performance may be seen when disks and DLT are accessed simultaneously (as during STORE/RESTORE). (Please read additional information in the Product Overview section for details on device performance.)

Unsupported connection Option #2

The HP-HSC F/W DA A2969A is not supported on HP 3000 S900 systems. This card requires the existence of "Device Adapter" Software that has yet to be written for MPE/iX.

Physical connection of these devices is accomplished via Fast/Wide/Differential SCSI Cables such as C2911A, C2924A, C2925A, C2905A, C2911A, C2924A etc.

NOTE While DLT Libraries are available for purchase, support for these devices is very limited on HP 3000 S900 systems. There are currently no plans for supporting the DLT Libraries for normal MPE/iX operations such as BOOT, FCOPY or STORE. DLT Libraries will be supported only when used in conjunction with a third party software package, Legato's NetWorker. Inquiries about DLT Library support should be made directly to Legato.

Product Overview

Product:	DLT8000
Standard, Non-compressed Capacity with DLT IV:	40 Gbyte
Average Maximum Data Transfer Rate To Tape:	~6 Mbyte/sec
Average Maximum Data Transfer Rate On SCSI Bus:	~10 Mbyte/sec
Maximum Devices Per Bus:	1

The DLT8000 drive accepts the new CompacTape IV cartridge. When this cartridge is used the amount of data you can store on a tape can be up to 40.0 GB native, or an expected 80.0 GB of data using the compression mode for the DLT8000. The actual amount of data stored in compression mode is dependent on the nature of the data being stored; i.e., text files tend to be highly compressible while code files are not. The device can also use the CompacTape III cartridge. The amount of data you can store on this tape cartridge can be 10.0 GB native capacity or 20.0 GB compressed.

Configuration

Here is an sample configuration as it should be on the HP 3000 for the DLT8000:

```
io> ld 91
LDEV:    91  DEVNAME:                OUTDEV:    0  MODE:
ID: DLT8000                RSIZE:    128  DEVTYPE: TAPE
PATH: 4/4.3.0                MPETYPE:   24  MPESUBTYPE: 8
CLASS: TAPE          DLTTAPE  DDUMP

io> lp 4/4.3.0
PATH: 4/4.3.0                LDEV:    90
ID: DLT8000                TYPE:    TAPE
PMGR: SCSI_TAPE2_DM        PMGRPRI:  10
LMGR: LOGICAL_DEVICE_MANAGER  MAXIOS:   0

io> lp 4/4.3
PATH: 4/4.3                LDEV:
ID: PSEUDO                TYPE:    DA
PMGR: TRANSPARENT_MGR    PMGRPRI:  6
LMGR:                MAXIOS:   0

io> lp 4/3
PATH: 4/3                LDEV:
ID: HP28696A                TYPE:    DA
PMGR: SCSI_DAM            PMGRPRI:  6
```

LMGR :

MAXIOS: 0

TERMDSM: User Interface Change under System Diagnostics

by Umakanth Rajasekhar
Commercial Systems Division

The underlying Diagnostics Interface has changed. This article gives information on what the user will see and what it means to invoke TERMDSM to diagnose the DTCs and its connected ports.

Starting with MPE/iX Release 6.5, a new diagnostic interface named 'Support Tools Manager' (STM) replaces SYSDIAG (which runs on SHERLOCK Diagnostics.) on HP 3000s. Most of the tools supported previously under SYSDIAG are still supported but they all now run under the MESA Diagnostic platform instead. Some SHERLOCK diagnostics have been replaced with native MESA tools. This article discusses the TERMDSM, the tool that diagnoses DTCs and its connected ports including X.25 connections.

What It Means to the User

Under MESA, Sherlock tools were migrated to MESA as expert tools and utilities. Only Expert tools requires a licence to run under MESA.

All the Diagnostic programs under SYSDIAG are now expert tools in MESA and utilities that run under SYSDIAG are still utilities. They do not require a license to run as before. TERMDSM still runs as a utility under MESA.

The functionality of TERMDSM in diagnosing DTCs and related ports is not changed; only the interface is changed.

What the User Will See

To invoke the utility, TERMDSM, the user has to run 'cstm' at the CI prompt. The user has to make sure that the 'diagmond' process is running in the system. The user will see a prompt from the interface (just like DUI in case of SYSDIAG). At that prompt the user must enter

```
ru termdsm
```

ru stands for Run Utility. Adequate help is provided with the interface and accelerator keys are also provided to lower the burden on the user to type in long commands.

A sample output from invoking TERMDSM is shown below:

```
MPE/iX: cstm
:cstm
Running Command File (/usr/sbin/stm/ui/config/.stmrc).

-- Information --
Support Tools Manager
Version A.11.00

Product Number B4708AA

(C) Copyright Hewlett Packard Co. 1995-1998
All Rights Reserved
```

Use of this program is subject to the licensing restrictions described in "Help-->On Version". HP shall not be liable for any damages resulting from misuse or unauthorized use of this program.

cstm>

cstm> ru termdsm

-- Termdsm Utility. --

```
*****
*****
*****          TERMDSM          *****
*****
*****          (C) Copyright Hewlett Packard Co.          *****
*****          1987, 1990, 1993, 1994          *****
*****          All Rights Reserved.          *****
*****
*****          Version C.01.60          *****
*****
*****  HP shall not be liable for any damages resulting from *****
*****  the use of this program.          *****
*****
*****          Catalog AM.060          *****
*****
*****
```

Welcome, Today is THU, AUG 19, 1999 at 4:02 PM

-- Utility : Section selection for TERMDSM. --

Enable Dialog For Additional Run Options? (y/n)[no]

Your system has these Distributed Terminal Controllers:

DTC Number	Network Address	Name
^		

Use this number as the "DTC number" in Termdsm commands.

Connect DIag DTc DUmp EEprom Help Reset Status Trace EXit ?

One addition to this menu driven tool is that MESA has more options to run the tool. For now, this is not suggested for TERMDSM.

In the above query which appears right after the banner showing TERMDSM,

-- Utility : Section selection for TERMDSM. --

Enable Dialog For Additional Run Options? (y/n)[no].

It is suggested to user to say no in the above to directly proceed to the tool prompt since there are no additional run options provided.

When the TERMDSM's prompt appears, it runs the same way as it used to run under SYSDIAG.

CAUTION The user is advised NOT to try reset/dump operation from a session connected to the DTC that he/she is diagnosing. This may potentially cause TERMDSM to have an abnormal exit.

Other Commands of Interest

Tools Related

UAL Utility Activity Log: This gives the status of invoking the tool.

UINF Utility Information

ABU Abort Utility: This aborts the tool from custom

Systems Related

In MESA, the user can run diagnostics for remote systems.

CS Connectsys: Adds systems and able to select the system to test.

DS Disconnectsys: Disconnects the (remote) system selected for diagnostics.

HP SureStore E Disk Array 12H on the HP 3000

*Rich Bassett
Walt McCullough
Commercial Systems Division*

Site, Software and Hardware Requirements

The following is a summary of Site, Software and Hardware requirements for support of the HP SureStore E Disk Array 12H.

Device installation and support is to be provided by HP. For detailed information on the support and configuration of the Model 12H AutoRAID, field support personnel should consult the updated version of the following documents:

- HP SureStore E Disk Array 12H -- *System Administrators Guide for HP-UX, MPE/iX and Microsoft Windows NT*
- *Users and Service Manual*
- *Control Panel Quick Reference Card*

These documents are located on the following external web page:

http://www.hp.com/essd/model12H_alldocs.html

NOTE The ARMserver software will be provided as a patch for MPE/iX 6.5.

AutoRAID Configuration Requirements

- Model 12H AutoRAID is supported using firmware revision 56
- Controller Mode should be set to Normal
- X Controller to be connected to the HP 3000, Y controller unconnected and terminated
- Rebuild priority should be set to High
- Logical Unit Number (LUN) (0) must be configured

Host Hardware/Firmware Requirements

- Support will only be allowed on a dedicated SCSI bus, with no other types of devices (homogeneous bus environment).
- Daisy chaining of two or more model 12H is supported within the recommended SCSI configuration.

NOTE Recommended SCSI Configuration: Maximum of 15 LDEVs per SCSI bus (Maximum of 8 LDEVs for performance.) Having 9 or more LDEVs configured will cause performance degeneration under high I/O loads.

- No support for shared SCSI busses.
- F/W SCSI Device Adapter Card (28696A, known as “Wizard) firmware MUST be at level 3728 or greater
- 9x9 Core I/O card (A2372-60004 or A3453-60010) will be conditionally supported with the following requirements:
 - Must disconnect all internal F/W SCSI drives
 - Must have firmware level of 3636.
 - Up to 8 LDEVs will be supported (Tested limit).
- Ldev 1 boot will have the following conditions:
 - Ldev 1 capacity limited to 4 Gbytes

Host Software Requirements

The Model 12H AutoRAID is supported on the MPE/iX 6.5 operating system with the addition of two software patches:

- MPE/iX C.65.01 (6.5)
- Software Patches:
 - The first patch contains the new `TYPE` manager and changes to `SDARRAY` for support of the new SCSI passthru functionality. This functionality is needed for the AutoRAID Management (ARM) utilities.
 - The second patch contains the ARM utilities. This patch is installed under the MPE/iX POSIX environment. Information on how to install this patch, see the *MPE/iX and SureStore E Disk Array 12H AutoRAID Manager Installation Guide*

NOTE The latest patches are available from the HP Patch web site.

- The `SYSGEN` product ID is `HPDARRAY`. See *Configuring AutoRAID SCSI LUNs for MPE/iX* within chapter 5 of the *System Administrators Guide for HP-UX, MPE/iX, and Microsoft Windows NT* for how to perform the `SYSGEN` configuration.

Host Diagnostic Support

Sysmap will identify the Model 12H AutoRAID as unsupported. For diagnostic support, the ARM utilities will replace the current `SHERLOCK` diagnostic. Note that `SCSIDSK2` can still be used to perform disk trouble analysis, but this is not a supported tool for the Model 12H AutoRAID product.

Off-line Diagnostics

There is a new off-line utility called `ARDIAG` that is supported within the support media off-line (ODE) diagnostics. See chapter 7 of the *HP SureStore E Disk Array 12H System Administrators Guide for HP-UX, MPE/iX, and Microsoft Windows NT* for how to use this command.

Configuring the 12H for Performance on the HP 3000

- The single X controller is the only connection to the HP 3000 supported for this release. It is recommended that you purchase the second Y controller in case of controller failure. The second controller will be able to take over when swapped into place, for it maintains a copy of the current firmware and disk maps
- If your 12H contains 12 18Gb disk drives we recommend configuring each of the LUNs for 11Gb of storage. This has shown to give the best performance versus capacity. This capacity allows for RAID1 and available space for two hot spare disks.
Total capacity of the LUNs should not exceed 90Gb. If this capacity is exceeded then you will trade performance with capacity.
- Data Resiliency Mode should be set to Normal. Some users have reported slow times for large restores. This problem can be corrected by setting the Data Resiliency Mode to Performance Mode only during the restore. Remember to reset the mode back to Normal when your large restore is completed.
- Rebuild Priority should be set to High. This may cause some degradation in OLTP but will reduce the likelihood of encountering a disk mechanical failure while the system is still recovering from a previous disk mechanical failure.

DVD Supported on the HP 3000

*by Alex Early and Dave Black
Commercial Systems Division*

The embedded CD-ROM shipped with many HP 3000 servers has come to the end of its life. However, many customers purchased HP 3000s with CD-ROMs to take advantage of system update material and manuals that are available on CDs. The A3715A, a 12x CR-ROM, is being replaced by the new DVD-ROM. The DVD-ROM supports usage of the older CD-ROM disks.

Customers wishing to take advantage of DVD-ROM for updates, can simply add a DVD-ROM to their HP 3000 system. For customers in support contracts, who have older CD-ROMs, your CD-ROM will be replaced with a DVD-ROM, if there are no CD-ROMs available. The DVD-ROM drive unit has a 32X CD-ROM transfer rate. The DVD-ROM is supported in the following packages:

Product Numbers

0950-3398	DVD-303 SCSI Slot-Load DVD-ROM Drive Unit
A5220A	Embedded Product Number.
C4318SZ opt 108	Highlander Factory Rack
C4314A	Highlander Deskside Standalone
C4315A	Highlander Field Rack

DLT7000 and Legato Storage Node Usage

Legato Storage Node and DLT libraries are now supported beginning with MPE/iX 5.5 with express 7 and MPE/iX 6.0 with express 1 plus patches. With Legato Storage Node software (B6266AA), the HP 3000 supports FWD SCSI connected DLT libraries with DLT7000s or DLT8000s directly connect to the HP 3000.

DLT7000 support is beginning with MPE/iX 5.5 Express 7 and MPE/iX 6.0 Express 1, no additional patches are required. DLT8000 require patch MPEKXU0A FOR MPE/iX 5.5 Express 7 AND MPEKXU0B FOR MPE/iX 6.0 Express 1. Legato Storage Node is supported in MPE/iX Express 7 and MPE/iX 6.0 Express 1 with no additional patches required.

DLT7000/8000 are also supported in MPE/iX 6.5 with no patches.

The following table provides a list of the Tape Libraries and their part numbers.

Product # Arrays	Description	Release	Patch	Date Available	Comments
A5583A	HP SureStore E Tape Library 2/20 (Rack)	MPE/iX 5.5 Exp 7+ MPE/iX 6.0 Exp 1+	MPEKXU0	Oct-99	HP 2 drives 20 slot DLT 8000 Library
A5584A	HP SureStore E Tape Library 2/20 (Desk)	MPE/iX 5.5 Exp 7+ MPE/iX 6.0 Exp 1+	MPEKXU0	Oct-99	HP 2 drives 20 slot DLT 8000 Library
A5585A	HP SureStore E Tape Library 4/40 (Rack)	MPE/iX 5.5 Exp 7+ MPE/iX 6.0 Exp 1+	MPEKXU0	Oct-99	HP 4 drives 40 slot DLT 8000 Library
A5586A	HP SureStore E Tape Library 4/40 (Desk)	MPE/iX 5.5 Exp 7+ MPE/iX 6.0 Exp 1+	MPEKXU0	Oct-99	HP 4 drives 40 slot DLT 8000 Library
A5587A	HP SureStore E Tape Library 6/660 (Rack)	MPE/iX 5.5 Exp 7+ MPE/iX 6.0 Exp 1+	MPEKXU0	Oct-99	HP 6 drives 60 slot DLT 8000 Library
A5588A	HP SureStore E Tape Library 6/660 (Desk)	MPE/iX 5.5 Exp 7+ MPE/iX 6.0 Exp 1+	MPEKXU0	Oct-99	HP 6 drives 60 slot DLT 8000 Library
A5597A A5599A A5618A	HP SureStore E Tape Library 20/700 DLT8000 Drive for Library 20/700 DLT7000 Drive for Library 20/700	MPE/iX 5.5 Exp 7+ MPE/iX 6.0 Exp 1+	MPEKXU0	Oct-99	HP 20 drives 700 slot DLT 8000 Library

HP SureStore E Disk Array XP256

by Rich Bassett--CSY Lab
Alex Early--CSY Product Marketing

The Hewlett-Packard SureStore E Disk Array XP256 provides high-capacity, high-speed mass storage, with continuous data availability, ease of service, scalability, and connectivity. The disk array can be connected to multiple server systems. The HP 3000 supports the XP256 through Fast-wide SCSI connections. The disk array can have up to 32 SCSI ports. It can be directly connected to the HP 3000 or connected via the SCSI-Fibre Channel Router for distances up to 10km.

The disk array is the first RAID disk array to truly provide continuous data availability. It is designed for nonstop operation and continuous access to all user data. The disk array has no active single point of component failure. It is not expected to fail in anyway that would interrupt user access to data. The XP256 is easily configured on the HP 3000.

Announcing Support of the HP SureStore E Disk Array XP256

The following is a summary of site, software and hardware requirements for support of the HP SureStore E Disk Array XP256. Device installation and support is to be provided by HP. For detail information on the support and configuration of XP256, refer to the *Operating System Configuration Guide*. Support of XP256 product will have the following general restrictions:

- Support will only be allowed on a dedicated SCSI bus, with no other types of devices (homogeneous bus environment).
- Daisy-chaining of two or more XP256 units are not supported.
- Maximum of 15 LDEVs per SCSI bus (Maximum of 8 LDEVs for performance)

NOTE Having 9 or more LDEVs configured will cause performance degeneration under high I/O loads.

- No support for shared SCSI busses.
- F/W SCSI Device Adapter Card (28696A, known as "Wizard") firmware MUST be at level 3728 or greater.
- 9x9 Core I/O card (A2372-60004 or A3453-60010, known as "Mustang") will be conditionally supported with the following restrictions:
 - Must disconnect all internal F/W SCSI drives.
 - Must have firmware level of 3636.
 - Up to 8 LDEVs will be supported (Tested limit).
- LDEV 1 boot will have the following conditions:
 - ONLINE HOT replacement of FRUs is supported with the exception of microcode downloading, online memory replacements, disk adapter replacements.

- Only RAID 1 is support for LDEV 1.
- Microcode down-loading, online memory replacements and disk adapter replacements must be performed with no jobs or sessions logged on.
- LDEV 1 capacity limited to 4 Gbytes (due to NIO IODC limits).

Support on MPE/iX has the following restrictions:

- No logging to the host of any vendor unique error codes. These codes are used to identify any hardware component (FRU) that has failed. For example, the failure of a fan inside the device would be reported as a FRU error.
- No predictive support.
- SYSGEN product ID is “HPDARRAY”.
- XP256 microcode revision level B++(rev 36) or later is required. Host mode table should be set to 08.
- Target address support only (Multi- LUN support available on request)

An additional requirement for support is that each XP256 will have a telephone modem connected for remote access. Normally, the XP256 will call HP support when it detects any errors. For information on support, see the HP support plan.

SureStore E Disk Array XP256 Business Copy XP and/or Continuous Access XP in an HP 3000 Environment

by Walt McCullough
Commercial Systems Division

Product Descriptions

Business Copy is an XP256 feature that allows for copy on demand of selected disk volumes. These volumes should be grouped as User Volumes, an MPE/iX term. These copied disks can then be split from their masters and presented to another HP 3000 machine for off-loading operations such as backup or application testing.

Business Copy volumes are managed from the internal XP256 PC or from a PC configured as the Remote Control PC.

The software used to manage Business Copy volumes is not currently available to execute directly from the HP 3000 but requires the separate PC to function.

Continuous Access is the feature that allows the XP256 to continuously maintain a copy of a selected number of disk volumes on another XP256 disk cabinet, located locally or remotely.

This feature is somewhat similar to Business Copy in that off-loading backup operations or application testing can be performed. The added value is that the user can use this feature to maintain a complete working set of data that could be used on the remote computer in the event of a catastrophic site outage at the primary location.

The software used to manage Continuous Access volumes is not currently available to execute directly from the HP 3000 but requires the separate PC to function.

Environment

Business Copy and Continuous Access is supported on MPE/iX release 5.5 with PPT07 and patch MPEKXL9 and MPEKXT8. It is also supported on MPE/iX release 6.0 with patch MPEKXL9 and MPEKXT8. For more information on patches and versions, please contact your support engineer.

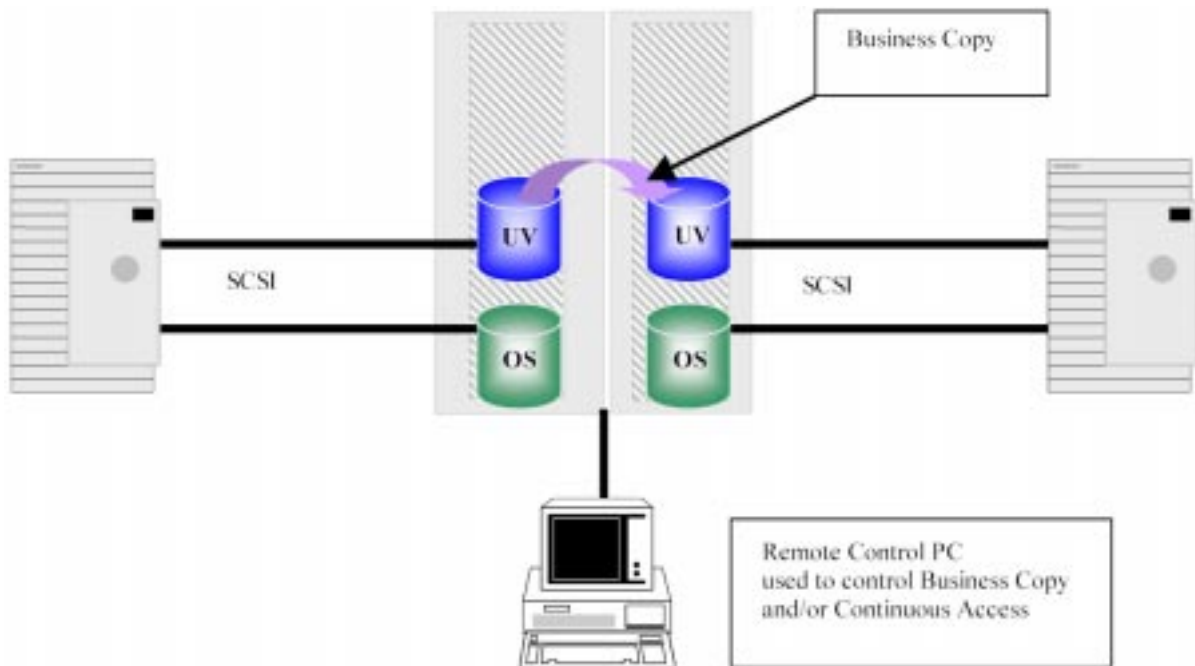
Basic Operation

Business Copy and Continuous Access have virtually the same procedures for setup. System A will “own” data located on a primary user volume set. A running copy of this data will be maintained on a Business Copy of the volume set.

At a customer defined time, the user will VSCLOSE the primary user volume set and then split the primary user volume set from the Business Copy volume set.

After the split is complete, the user can then VSOPEN the primary user volume set onto System A and continue operations. System B can now VSOPEN the split off Business Copy volume set and proceed with its operations.

Figure 4-1. Basic Setup of Business Copy



System A Procedures

Setup System A

The first time setup requirements are as follows:

1. Create the Logical Unit Numbers (LUNs) and paths you will use on the XP256 and then using Sysgen or IOCONFIG create the LDEVs associated with the paths and LUNs. Make sure to use HPDARRAY as the ID type. We use the LDEV 60 through 62 for our examples.

Test that MPE/iX recognizes the new LDEVs by either rebooting and running ODE Mapper or if you used IOCONFIG to issue the commands, use `DSTAT ALL`.

```
(MPE/iX Prompt):dstat all
LDEV-TYPE   STATUS      VOLUME (VOLUME SET - GEN)
-----
60-OPEN-3   UNKNOWN
61-OPEN-3   UNKNOWN
62-OPEN-3   UNKNOWN
```

(MPE/iX Prompt):

2. Now create the User Volume Set that will be used to off-load processing. The volume set name for this example is `PRODUCTION_DATA`. Use the `VOLUTIL` utility program to create the user volume set named `PRODUCTION_DATA`. Once the volume set is created, `VSCLOSE` and then `VSOPEN` the volume set before proceeding onward.
3. Issue the `DSTAT ALL` command and the following should appear:

```
(MPE/iX Prompt):dstat all
  LDEV-TYPE      STATUS      VOLUME (VOLUME SET - GEN)
-----
60-OPEN-3      MASTER      MEMBER1      ( PRODUCTION_DATA-0 )
61-OPEN-3      MEMBER      MEMBER2      ( PRODUCTION_DATA-0 )
62-OPEN-3      MEMBER      MEMBER3      ( PRODUCTION_DATA-0 )
```

4. Next, create the directory and accounting structure on the system that will use the User Volume. Be sure to remember to use the proper syntax and parameters *ONVS* and *HOMEVS* when creating the groups. You are now ready to restore or create the files that will reside on the user volume set *PRODUCTION_DATA*.

From the Remote Control PC, configure the LUNs and paths for the Business Copy volumes and start the copy operation. This takes about 1 GB per minute depending on your configuration. Once this operation has completed, you may proceed.

At this point you should have a system configured with a user volume set name *PRODUCTION_DATA* and a directory and accounting structure that is accessible from System B. You should also have a copy of this volume set maintained by Business Copy. To access the files and directories when this volume set is “moved” over to System B, you need to have the identical directory and accounting structure on System B that is already on System A.

BULDACCT To access those files on the Business Copy volumes you must first create the accounts and groups on System B, that reference the directory structure located on the user volume. Re-issue the *NEWACCT* and *NEWGROUP* commands on System B or use *BULDACCT* to create a job stream that will create the directory for you.

5. Execute the *buldacct* utility program and for this example three accounts exist on the *PRODUCTION_DATA* volume set. *SALESDB*, *ORDERDB* and *STOCKDB*.

```
:RUN BULDACCT.PUB.SYS;INFO="SALESDB,ORDERDB,STOCKDB%VSACCT=PRODUCTION_DATA"
```

This will create 2 files *BULDJOB1* and *BULDJOB2*. These files create the directory structure and re-create the UDC linkages. Store these files off and restore them to System B. (*DSCOPY* can also be used)

System B Procedures

1. Either re-issue the *NEWACCT* commands that pertain to creating the accounts on the system volume set and only those *NEWGROUP* commands that you used the parameter *HOMEVS* or restore the *BULDJOB1* and *BULDJOB2* files onto System B and stream those jobs.
2. From *Sysgen* or *IOCONFIG* issue the commands to configure the *LDEVs* of the Business Copy volumes. Be sure that the copy has completed.

If you used *Sysgen* to configure the volumes you will need to reboot the system. If *IOCONFIG* was used, then it will try to mount the volumes as you add the *LDEVs*.

Once the volumes are configured and the directory structures are intact, issue the *dstat all* command.

```
(MPE/iX Prompt):dstat all
  LDEV-TYPE      STATUS      VOLUME (VOLUME SET - GEN)
```

Volume	Attributes	Member	Volume Set
60-OPEN-3	LONER-RO	MEMBER1	(PRODUCTION_DATA-0)
61-OPEN-3	LONER	MEMBER2	(PRODUCTION_DATA-0)
62-OPEN-3	LONER	MEMBER3	(PRODUCTION_DATA-0)

The master volume should say LONER-RO, which means that it is closed and access is defined as Read-Only. MPE/iX will not let you mount a Read-Only disk because it is not supported on MPE/iX. The Read-Only attribute is set by the XP256 while it is active copy mode.

Business Copy Procedures

On System A

```
VSCLOSE PRODUCTION_DATA;NOW
```

This will quiesce I/Os to the volume set and flush all files and data structures resident to that volume set to disk.

Within the Remote Control PC make sure that the primary and Business Copy disks are in sync and then split the primary volumes from the Business Copy volumes.

(After the split you may VSOPEN the PRODUCTION_DATA volume set on System A and continue operations)

On System B

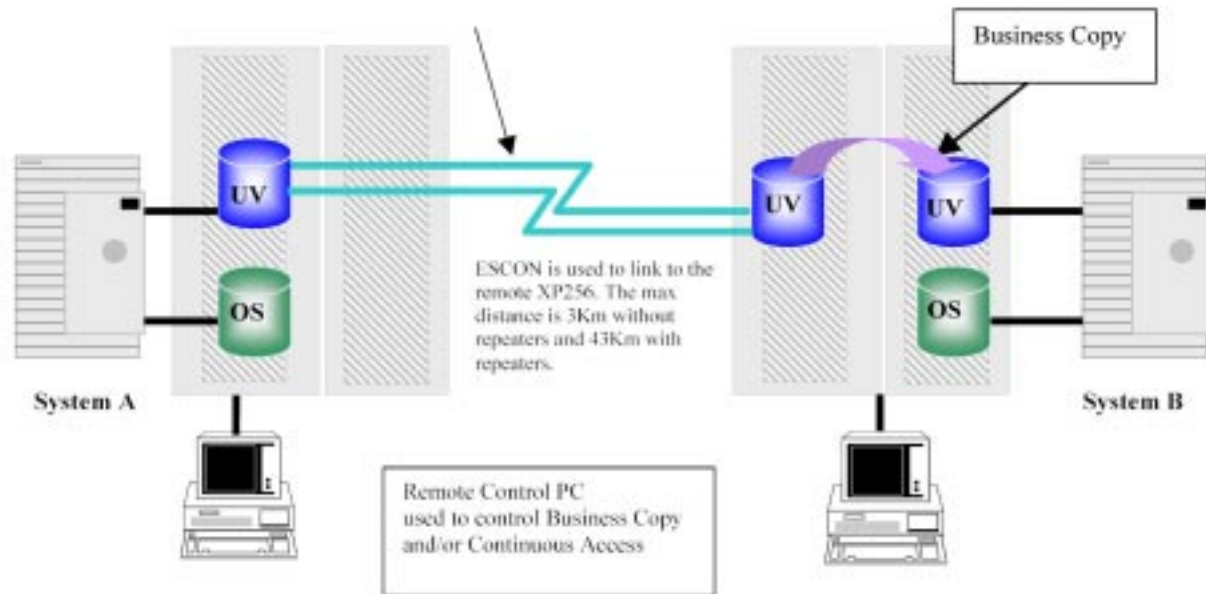
After the primary volumes have been split from the Business Copy you can issue the VSOPEN PRODUCTION_DATA command on System B and start operations.

When operations are completed, VSCLOSE PRODUCTION_DATA;NOW and from the Remote Control PC re-establish the primary and Business Copy mode. There is no need to log off users from System A to re-establish Business Copy.

Continuous Access

The diagram below describes the use of both Continuous Access and Business Copy to provide both disaster recovery and off loading of work done by the remote data center.

Figure 4-2. Using Both Business Copy and Continuous Access



The goal is to always maintain a link to the primary volume set through Continuous Access and use a separate set of disks that are maintained by Business Copy for planned events like work off-loading.

The process is almost the same as that of Business Copy. First quiesce I/O and log users off the `PRODUCTION_DATA` volume set by issuing the `VSCLOSE;NOW` command on System A. Then when all data has been moved to the remote XP256 only then do you split the Business Copy volumes from the Continuous Access volumes on the remote XP256.

Only when the Business Copy volumes are split on the remote XP256 can System B `VSOPEN` its copy of the `PRODUCTION_DATA` volume set.

5 Internet and Interoperability

This chapter provides the following articles about the Internet and interoperability.

- Large Transactions for IMAGE Users
- XM User Checkpoint Priority Control
- Transact/iX Expanded B-Tree and File Open Enhancement
- Apache for MPE/iX
- Introducing LDAP C-SDK/iX
- MPE/iX Software Developer's Kit (SDK) for Java, Version 1.2
- HP Driver for JDBC

Large Transactions for IMAGE Users

*By B T Vikram Kumar, B S Jyoti, Shobha Pradeep
Commercial Systems Division*

Background

Of late many of the users of TurboIMAGE and IMAGE/SQL have the requirement of transaction sizes larger than is currently possible today (4 MB). This is true, especially for users who run applications involving bulk puts/ deletes /updates. The situation becomes more difficult with the limitation of the size of the Transaction Manager (XM) userlog, if the transactions are on the same volume set. This integrated solution involving TurboIMAGE/iX, XM, IMAGE/SQL and ALLBASE/SQL addresses these problems. The following sections discuss these issues, and various options available to the users to overcome these issues, as a result of this solution.

Current Limitations

Version C.03.00 of TurboImage introduced the concept of dynamic transaction, meaning the application can rollback a transaction at runtime. Later when the Transaction Manager started handling the dynamic transactions, some customers used to see the stalled transaction message due to the limit on the size of the transaction.

Transaction Size Dynamic transactions are the ones which are bracketed by `DBXBEGIN` and `DBXEND` or `DBXUNDO`, which may span a single database or multiple databases. In the case of multiple database transactions the total number of databases which can take part in a transaction is restricted to 15. A dynamic transaction can be rolled back using `DBXUNDO` or automatically when the application aborts within a transaction. When an application does a bulk insert, delete or update, there is a chance that it hits the current limit of 4 MB on transaction size. Until recently, the process used to abort under such situation, after displaying an error message that the transaction size has exceeded. An intermediate solution has been provided (patch `MPEKX35`) to rollback the transaction before process abort, so as to leave the database in a consistent state, and user could re-run the application after manually pruning the application. However, this was not a comfortable alternative, and users have been requesting that there be an increase in the transaction size.

User Log Size The user log related limit arises due to the fact that XM in MPE is log-based. That is all the transactions, mainly the ones which involve manipulation to the database, are logged on to a circular log file. When there are many applications writing into the XM log file, it could also get filled up easily, again resulting in process abort. The userlog (all references to userlog in this article refers to XM userlog) comes into the picture here because all the changes to IMAGE databases on a specific volume set are written to the same userlog. If only the transaction size limit is increased, because of the increased number of open transactions, the 64 MB limit is easily attained.

Solution Objective

The main objective of this solution therefore has been to increase the transaction size from the existing 4 MB. The second objective is to increase the size of the XM userlog from the current value of 64 MB to support large transactions. The third objective is to ensure compatibility with existing applications, so that they can run without any problems even with these enhancements.

Overview of the Solution

In order to achieve the above mentioned objectives, the capabilities of the following products /subsystems of MPE/iX have been enhanced:

- Transaction Manager
- TurboImage/iX
- Image/SQL
- ALLBASE/SQL

Transaction Manager

A patch MPEKXG8 was released (on MPE/iX 5.5) to cater to the increase in number of transactions by reducing the data stored in a transaction.

This was made possible, by keeping only the essential information required for possible rollback in future. A large amount of internal data becomes irrelevant after completion of an extended transaction (we will refer any dbput/ dbdelete/ dbupdate as an extended transaction throughout this article). This allows for more transactions even within the current 4 MB limit.

However, the long term solution addressed here, calls for increasing the transaction limit to larger values. Based on current OS limits, the maximum transaction size has been increased to 32 MB in this solution. In addition to this, XM has been changed to send out an early warning message (soft limit warning), before the transaction size limit of 32 MB is hit (i.e. at 28 MB). When, ultimately the transaction size reaches 32 MB, it aborts the process and rollback the transaction.

Another enhancement in XM is to increase the userlog. The current userlog size is 64 MB per volume set. Currently, for each transaction, the information that can be contained in the userlog is limited to 4 MB. Due to the increase in the transaction size for each process (to 32 MB), the total userlog file should also be increased. The userlog resides on the master volume of a volume set. User can use VOLUTIL utility to increase the size of the userlog (the default value is 64 MB) with the following,

```
>From VOLUTIL,  
ALTUSERLOG <volset-name> <new size>
```

This command accepts the volume set name and the new size of XM userlog for it in sectors. It can not be less than 264256 sectors corresponding to 64 MB. One caution is that ldev 1 has limited reserved space and one can not expand the XM userlog beyond that.

The actual expansion of the XM userlog happens when user boots the system next time. This is because the XM log files can not be expanded/truncated while in use.

During bootup, just after recovery of the userlog, the XM userlog is expanded and a

message is written onto the console.

One can also truncate the userlog after expansion if there are any disk space or performance issues. The `volutil` interface is same for this operation too.

One can backdate the OS to an earlier version with the userlog in the expanded state.

The expansion of the userlog may fail due to one or more reasons. In that case, the system reports it through a warning message on the console.

If it fails due to lack of disk space, make sure that there is sufficient contiguous space on the master that would fit the need for the increment in the userlog size. Then reboot the system so as to enable it to complete the expansion of userlog. You must ensure that memory beyond 256 MB is available on the system, before increasing the userlog through `VOLUTIL`.

NOTE The userlog size gets scaled between 32 MB and 64 MB if the memory size available on the system is > 64 MB and < 256 MB.

The current log size can be seen through the command `SHOWUSERLOG`.

```
>From VOLUTIL,  
SHOWUSERLOG <volset-name>
```

This command accepts a valid MPE volume set name and shows the current size of the XM userlog in sectors.

It has to be noted that, a single transaction can grow up to 32 MB now and if the size of userlog is 64 MB, total size of all active transactions on that volume set can be 32 MB (half log size). If the total size of all active transaction exceeds half log size, all the processes which are actively logging into the same userlog are aborted. So, a large transaction may result in all processes getting aborted, if the userlog size is not expanded appropriately.

The downside of having a larger userlog is that the XM checkpoint process has more data to write to disk at reaching the end of each userlog half. This can have effect on response time to on-line users.

TurboIMAGE

Support for Large Transaction Size: Image has been enhanced to support the new 32 MB transaction size provided by XM. This means, without doing anything special, one can have approximately more than 8 times the number of extended transactions in a dynamic transaction than possible today (including the XM optimization available with XM patch `MPEKXG8`, which has been rolled into this solution).

User Control of Transactions: In addition to the benefit of transaction sizes up to 32 MB, available across the board to all IMAGE and IMAGE/SQL users, direct IMAGE applications benefit from a new feature. This feature helps a user to judiciously control/manage large transactions. User application can avail of an early warning message feature through which the application can have virtually a full control over the transaction. To avail this facility, a new mode (mode 18) has been added to `DBCONTROL` intrinsic. If the application makes a call to `DBCONTROL` mode 18 before the start of a transaction, it sets a flag in IMAGE, to trap the `softlimit` warning message from XM, and inhibits all new transactions. Now the user application have two options:

- to commit the transaction (through a call to DBXEND)
- to rollback the transaction (through a call to DBXUNDO)

Thus, there will not be a situation in which the transaction size reaches the hard limit resulting in automatic rollback and process abort.

Once the decision on rollback/ commit has been made and executed, you can start a new dynamic transaction to continue the job until the next softlimit warning is reached, and respond as above. The softlimit has been chosen to be at 28 MB. At any stage, if it is desired not to use this feature, the new DBCONTROL mode-19 can be used to reset the flag in IMAGE, so that no early warning, or forced commit/rollback occurs. To have compatibility with existing applications, the flag is not set by default, that is, by default there won't be any user control on the transaction size, which can grow to 32 MB.

Once 32 MB has been reached, XMI aborts the process with a "Stalled Transaction" message, and IMAGE automatically rolls back the transaction and leaves the database in a consistent state.

A typical usage sequence to make use of the new feature is:

- DBCONTROL Mode 18
- DBXBEGIN
- Intrinsic like DBPUT, DBDELETE or DBUPDATE
- Receives soft limit warning (message # -1025) from the modify intrinsic DBPUT/DBDELETE/DBUPDATE, when softlimit is reached.
- DBXUNDO / DBXEND

NOTE DBCONTROL Mode 18, has to be called before the start of a dynamic transaction, that is, before DBXBEGIN. (Otherwise, it returns an error message, and does not have any effect on the transaction manager). However, DBCONTROL Mode 19, can be called anywhere in the program. It turns off the detection of the softlimit warning message from XM, regardless of whether it is set or reset.

Image/SQL and ALLBASE/SQL

For the benefit of users doing bulk deletes, ALLBASE has a "DELETE WITH AUTOCOMMIT" feature for ALLBASE tables. Now this facility has been extended to IMAGE tables as well. IMAGE/SQL and ALLBASE/SQL have been enhanced to support this. The command would be "DELETE WITH AUTOCOMMIT FROM" owner.table [WHERE search conditions].

However, this option has to be exercised with great care, as IMAGE/SQL does periodic commits at internally maintained intervals, and there is not any way to rollback the deletes. For bulk deletes, the autocommit feature is very useful in that, chances of hitting the 32 MB limit is remote.

Limitations

If MPE/IX 6.5 is installed on a system, then TurboIMAGE/iX version C.08.00 also needs to be installed. However, this Image version can work with older versions of XM as well. In this case, XM will support transaction sizes only up to 4 MB, and will not issue a softlimit warning message.

If there are several processes having large transactions, and if XM user log has not been expanded appropriately, all the processes may get aborted once the log becomes full.

Summary of Benefits to the Users

- Across the board benefit of transaction sizes up to 32 MB for all IMAGE and IMAGE/SQL users
- Can accommodate more number of extended transactions in a top level dynamic transaction
- Can accommodate more number of extended transactions in a top level dynamic transaction
- No need to change existing applications to get the benefit of increased transaction size
- User control of dynamic transaction, if desired, to avoid stalled transaction
- Facility of "DELETE WITH AUTCOMMIT" feature to IMAGE tables
- User configurable XM userlog

XM User Checkpoint Priority Control

Jyothi BS
Commercial Systems Division

Background

Whenever one half of the Transaction Manager (XM) user log gets full, XM starts the checkpoint process to write out all the dirty pages of the files which have transactions logged in that half.

During user Checkpoint on many busy high-end systems, the response to on-line users is very poor. Sometimes, it almost stalls. The main reason is that there is too much of disk i/o being used at very high priority, i.e., C152. This problem may increase further if the userlog is expanded from the current 64 MB.

Solution Objective

The user should have the flexibility to easily alter the XM user chkpt priority (at user level). The user-given priority value should be persistent across system boots.

Overview of the Solution

The goal of this solution is to be able to reduce the impact of XM user checkpoint operation on the on-line response. This means controlling the priority at which chkpt i/o's are issued by MM, and also having the priority configurable by the user.

The priority is configurable at volume-set level and it is not system wide. The priority for the required volume-set can be changed using `volutil`.

How and When to Control the Checkpoint Priority

The users should lower the priority only if they see a significant impact on-line response time. It should not be made so low (such as 255) that checkpointing takes a long time and system nearly stalls. It should not be so high (such as 152) that during checkpointing, on-line response comes down drastically.

How to Give the Request to Change the Checkpoint Priority

The `volutil` tool is enhanced to provide two new commands to facilitate the control of `chkpt` priority. They are,

```
SHOWCHKPTPRI <volset-name>
```

This command displays the priority of the checkpoint process for the given volume set.

```
ALTERCHKPTPRI <volset-name> <new priority>
```

This command modifies the priority of the User checkpoint process for the given volume set. The checkpoint priority should be in the range of 152 to 255, with the recommended value at 200. Priority 0 resets the checkpoint priority to default which is C152. The changed priority becomes effective at the next checkpoint.

Transact/iX Expanded B-Tree and File Open Enhancement

by James Overman
Commercial Systems Division

Hewlett-Packard has implemented two more enhancements in Transact/iX (HP30138A.07.00). They are:

- Expanded support of TurboIMAGE B-tree
- Ability to pass database and file ids to Transact/iX

Support of B-Trees in Transact

With the release of IMAGE/SQL that provided TurboIMAGE B-trees, Transact/iX supported a portion of that feature. The Simple Argument structure of DBFIND was supported requiring no modifications to the Transact language. A FIND for X or U type items having B-tree indices and with the BTREEMODE1 option **ON** performs index searching with a wildcard character.

Here's an example. First, assume the BTREEMODE1 option is set **ON** for this database. Also, B-tree indices have been created for the key item of the master data set allowing for indexed retrievals. See the *TurboIMAGE Reference Manual* if you are not familiar with B-tree indexing.

In Transact, the source code would look like the following:

```
MOVE (NAME) = "Sm@" ;  
SET(KEY) LIST(NAME) ;  
OUTPUT(CHAIN) PEOPLE, LIST=(NAME) ;
```

This program would display all names in the detail data set, PEOPLE, which begin with "Sm".

```
Smalley  
Smith  
Smythe
```

The full capability of B-trees required more functionality to be added to Transact/iX. HP has expanded upon the Simple Argument DBFIND by allowing other DBFIND modes to be specified. By allowing the DBFIND mode and the argument (structured or simple) to be specified, Transact/iX supports the full capability of TurboIMAGE B-trees.

The Structured Argument can be specified by the **SET(KEY)** and **LIST(KEY)** statements. The **SET(KEY)** sets both the argument and key registers to the search argument item. The **LIST(KEY)** statement resets the key register to the search or key item. The Structured Argument item can be constructed via parent/child items or with simple items using string functions and concatenation.

Transact/iX has an added FINDMODE option on DELETE, FIND, GET, OUTPUT, REPLACE, and PATH statements. The syntax looks like:

```
FIND(CHAIN) dataset, LIST=(item1:item3), FINDMODE=4;
```

The B-tree search can only be used for the following data retrieval verbs: DELETE, FIND, GET, OUTPUT, REPLACE, and PATH. Only the CHAIN or RCHAIN modifiers can be used for the data retrieval verbs. The key and argument register should be set prior to the data retrieval. The STATUS count returned may not always be accurate. For example, DBFIND modes 21 and 24 do not return the correct chain count or entry numbers.

The complete syntax for the FIND verb is shown below.

Syntax

```
FIND [(modifier)] file-name [,option-list];
```

Parameters

<i>Option-list</i>	One or more of the following options, separated by commas. Only the additional option is listed here. See the <i>Transact/iX Reference Manual</i> for all the other possible options.
FINDMODE=mode	Only available with the CHAIN or RCHAIN modifiers. Performs optimized searching of an index by using TurboIMAGE/XL's B-tree index functionality or Third Party Indexing. The mode is the mode that is specified for the DBFIND of this data retrieval verb. If the FINDMODE option is not specified, DBFIND mode 1 is used.

Here are some examples of B-tree Structured Arguments in Transact/iX. Assume BTREEMODE1 is ON and B-tree indices exist.

Example 1

```
system ex1,base=passks;
define(item) search-arg      x(30):
                    search-type  x(2)=search-arg(1):
                    search-version I(5,,2)=search-arg(3):
                    arg1-size    I(5,,2)=search-arg(5):
                    arg2-size    I(5,,2)=search-arg(7):
                    arg1        x(8)=search-arg(9):
                    arg2        x(8)=search-arg(17):
                    prod-no      x(8):
                    ship-date    x(6);

list search-arg,init:
    prod-no:
    ship-date;

<< Retrieve all product numbers beginning with PROD. >>
<< Uses DBFIND mode 1. >>

move (prod-no) = "PROD@";
set(key) list(prod-no);
output(chain) inventory, list=(prod-no:ship-date);

<< Retrieve all product numbers between PROD1 and PROD2. >>
<< Uses DBFIND mode 4. >>
```

```
<< Uses parent/child items to devise Structured Argument. >>

move (search-type) = "[";
let (search-version) = 0;
let (arg1-size) = 8;
let (arg2-size) = 8;
move (arg1) = "PROD1 ";<< Blank pad the full number of characters.>>
move (arg2) = "PROD2 ";<< Blank pad the full number of characters.>>

set(key) list(search-arg);
list(key) prod-no;
output(chain) inventory, list=(prod-no:ship-date),findmode=4;
exit;
```

Example 2

```
system ex2, base=passks;

define(item) search-arg      x(30):
                search-version  I(5,,2):
                arg1-size      I(5,,2):
                arg2-size      I(5,,2):
                arg2           x(8):
                prod-no        x(8):
                ship-date      x(6);

list search-arg,init:
    search-version,init: << 0 >>
    arg1-size,init:
    arg2-size,init:
    arg2,init:
    prod-no:
    ship-date;

<< Use items and functions to concatenate the Structured Argument. >>
<< Use DBFIND mode 24. >>
<< Same output as the second part of Example 1. >>

<< search-version set to 0 during list statement >>
let (arg1-size) = 8;
let (arg2-size) = 8;
move (arg2) = "PROD2 ";<< Blank pad. >>

move (search-arg) = "[" + (search-version) + (arg1-size) +
                    (arg2-size) + col("PROD1",9) + col((arg2),17);

set(key) list(search-arg);
list(key) prod-no;
output(chain) inventory, list=(prod-no:ship-date),findmode=24;

<< Another way to setup the Structured Argument. >>
<< Same results as above. >>

list search-arg,init;
move (search-arg) = col("[",1);
move (search-arg) = col((search-version),3);
move (search-arg) = col((arg1-size),5);
move (search-arg) = col((arg2-size),7);
```

```

move (search-arg) = col("PROD1",9);
move (search-arg) = col("PROD2",17);

set(key) list(search-arg);
list(key) prod-no;
output(chain) inventory, list=(prod-no:ship-date),findmode=24;
exit;

```

File Open Transact

"File Open" is the name of the feature which allows database and file ids to be passed into Transact/iX from other languages or Third Party packages. This allows the Transact program to use the same databases and files used by the calling program. The Ids are shared so both programs can share current record numbers, entries, paths, etc. This feature requires a call to `TL_CALL_TRANSACT` (aka Transact/iX ACI) from another language or package. HP has also included an additional enhancement, the ability to defer the database open.

These enhancements required changes to two Transact/iX statements, `SYSTEM` and `LIST`.

SYSTEM Statement

The `SYSTEM` statement needs to tell Transact/iX if the database or file id is being passed. The new syntax is shown below.

Syntax

```
SYSTEM program-name [,definition-list];
```

Parameters

Definition-list Description of the files or data sets used during execution. See the *Transact Reference Manual* for other definitions.

```
BASE=base-name([["password"][, [mode][[, [optlock][[, [basetype][[, [open-type]]]]]]]).
```

open-type

Used by Transact/iX to determine when and if to open the database. The valid types are OPEN, DEFER, PASSED. OPEN is the default and opens the database at the beginning of the Transact/iX program. DEFER delays the database open until the database is first accessed. PASSED tells Transact/iX not to open the database - the database id is being passed from another program.

```
FILE=file-name([ [access][ (file-option-list) ] [, [record-length]
[, [blocking-factor] [, [file-size][[, [extents]
[, [initial-allocation][[, [file-code]]]]]]]]]).
```

file-option-list

Any of the following fields provided that they do not conflict in meaning: OLD, NEW, TEMP, \$STDLIST, \$NEWPASS, \$OLDPASS, \$STDIN, \$STDINX, \$NULL, ASCII, CCTL, SHARE, LOCK, NOFILE, HP3000_16, HP3000_32, DEFER, PASSED.

Either DEFER or PASSED can be specified. DEFER opens the

file when the file is first used in a Transact/iX program. PASSED does not open the file - generally the file id is passed from another program. If neither DEFER or PASSED is specified, then DEFER is assumed. DEFER is the way Transact/iX worked prior to this enhancement.

```
KSAM= file-name[(access)[(file-option-list)]].
```

See FILE= for DEFER and PASSED options.

LIST Statement

The next statement that Transact/iX needed to enhance is the LIST statement. The new syntax is shown below.

Syntax

```
LIST item-name[,option-list]; << Adds the item-name to the list >>  
    << register.>>
```

Parameters

<i>Option-list</i>	Specifies a value to be placed in the data register. For current options, refer to the <i>Transact Reference Manual</i> .
BASEID[(base-name)]	An X(2) item that contains the database id passed from a main program assigned to a database by TurboIMAGE's DBOPEN. If no base-name is specified, then the home base is assumed. Note, the home base cannot be specified.
FILEID(file-name)	A 16-bit integer (I(5,,2)) containing the file identifier assigned to a file by MPE's FOPEN.

The database/file on the SYSTEM statement should have the PASSED option or an error will occur. Any database/file name in the SYSTEM statement can actually be a different named database/file. There is not a check to verify that the database/file names are actually the same.

For example, if the calling program opened a database named PEOPLE, then Transact/iX assigns that database id to its database name of EMPLOYEES (structured the same as PEOPLE):

```
LIST people-id,BASEID(employees);
```

Example: Pascal Code - Main Program

Here is an example using Pascal to call Transact/iX.

```
.  
. .  
. .  
type  
  data_record = packed record  
    x_type   : packed array[1..8] of char;  
    id_type  : packed array[1..2] of char;  
    filler   : packed array[1..48] of char;
```



```

    end;
.
.
.
var
    data_buffer : data_record;
.
.
.
base_name := ' NAME; ';
password := 'PASS1 ';
mode := 1;
dbopen (base_name, password, mode, db_status_array);
{ check db_status_array }
.
.
.
strmove (2, base_name, 1, data_buffer.id_type, 1);

system_name := 'aci01  ';
data_size := 58;
tl_call_transact (system_name, addr(data_buffer),
                 data_size, return_status);
{ check Transact's return_status }
.
.
.
end.

```

Example: Transact Code - Subprogram

```

<< This code is compiled and added to the XL. >>

system aci01,base=name(,1,,PASSED); << NAME database is not opened >>
                                     << by Transact/iX. >>

define(item) x-item      x(8):
              name-id    x(2):
              name       x(8):
              address    x(32):
              phone      x(8);

list      x-item:
          name-id,BASEID:  << Home base assumed. >>
          name, init:
          address,init:
          phone,init;

<< Use database ID from Pascal code to perform retrieval. >>

output(serial) employees, list=(name:phone);
exit

```

Apache for MPE/iX 6.5

by *Barbara Dubbert*
Commercial Systems Division

Overview

Apache for MPE/iX is server software which turns your HP 3000 into a full-featured web server. With the Apache web server, HP 3000 users can now do business over the Internet.

As a web server, your HP 3000 can provide users with direct access to documents and applications residing on your system. These applications can include both internet and intranet dynamic database connectivity using a browser as a common, easily-maintained interface. The client browser can be any one of a variety of browsers including those from Microsoft and Netscape.

The Apache Web Server is currently the most popular web server on the market with over 50% of the web server market share. Apache on the HP 3000 joins a growing number of computer platforms which support Apache including HP-UX and other UNIX derivatives, Linux, Win32 and various other server and desktop platforms.

Feature Set

Apache for MPE/iX supports a rich set of features including:

- The latest HTTP protocol, HTTP/1.1
- Advanced Logging
- Access Control
- Common Gateway Interface applications (CGI)
- Server Side Includes (SSI)
- Cookies
- Imagemaps

For details on installing and using Apache for MPE/iX, refer to the Technical Overview section.

Technical Overview

As a web server, your HP 3000 can provide users with direct access to documents and applications residing on your system. Users make requests to the web server via a client browser using the Hypertext Transfer Protocol (HTTP). The sole purpose of a web server is to translate a request (URL) into either a filename, and then send that file back over the network, or into a program name, and then run that program and send its output back.

Once you start the executable, `HTTPD`, Apache runs silently in the background, waiting for a client's request to arrive on a port to which it is listening. Apache listens on the port specified in its configuration file.

When a request arrives, Apache hands the request to one of its child processes to service and returns to listen again on the port.

Major Components

The major components for Apache are the web server program (HTTPD), the job stream file which runs the HTTPD program (JHTTPD), a set of configuration files for enabling Apache features, the `htdocs` directory containing HTML pages, the `logs` directory, and the `cgi-bin` directory. Apache also comes with a full, on-line manual set.

- HTTPD is the Apache web server program. “HTTP” stands for the protocol used between the client browser and the Apache web server. “D” stands for daemon, a system program which automatically handles certain system operations. The HTTPD web server program intercepts an incoming request from the browser, interprets and handles the request, then delivers output in the form of an HTML page to the client. The HTTPD program file resides in the `/APACHE/PUB` directory.
- JHTTPD is a job stream file which runs the Apache web server as a standalone process. This file also specifies the system timezone and the location of the global Apache configuration file, `http.conf`. JHTTPD resides in the `/APACHE/PUB` directory.
- The configuration files tell the HTTPD server program where to find files, which files and directories can be accessed, who can access them, and the location of executable programs on the server. These files give you the flexibility of having the web server do only what you want it to do. These files reside in `/APACHE/PUB/conf`.
- The `htdocs` directory contains the public documents, images, and data to be served to clients. The `htdocs` directory and the directories below it are available to anyone accessing your web server.
- The `logs` directory contains log data of both web server accesses and of errors.
- The `cgi-bin` directory contains CGI scripts. These are executed by Apache on behalf of its clients.
- The `bin` directory contains a set of scripts and programs contributed by the Apache Software Foundation. They are included as examples only and are not supported by HP.
- The Apache manual set, residing in `/APACHE/PUB/htdocs/manual`, describes the entire Apache feature set and the syntax definition for configuring these features. Access to the manual documents is specified with the URL, `http://yourserver.com/manual/index.html`, where “yourserver.com” is the domain name of your HP 3000.

Steps for Starting Apache

To start the Apache web server:

1. Prepare your system for network access
2. Copy and edit the configuration files
3. Edit the JHTTPD job stream file
4. Start the HTTPD web server program
5. Verify that the HTTPD web server is running

Prepare Your HP 3000 for Network Access

Before your HP 3000 can act as a web server, it must be available for network access via TCP/IP. In preparation, you should:

- Configure TCP/IP on your system
- Have a domain name associated with your system's IP address

Apache communicates on the network using the HTTP Hypertext Transfer Protocol which, in turn, uses TCP/IP. Using `NMMGR`, configure your system's IP address and subnet mask. Then run `:NETCONTROL START` from the CI command line and verify that it ran successfully.

You will also want to get a domain name. This is a unique identifier such as "yourserver.com" which is used (instead of the IP address) to direct requests from a browser to your server. Request a domain name from the administrator of the Domain Name Server (DNS) on your network.

Configure Apache

Apache reads two global configuration files when it starts: `httpd.conf` and `mime.types`. These configuration files determine how Apache behaves. Earlier versions of Apache read two additional global configuration files: `access.conf` and `srml.conf`. These additional global configuration files can still be used, but by default they are empty and their original content is now included in the `httpd.conf` file.

Edit your Configuration Files

The `mime.types` file comes as `mime.types.default` and must be copied or renamed to `mime.types`. Similarly the `magic.default` file must also be copied or renamed. The `httpd.conf` file comes in two different versions. The `httpd.conf.default` file is for a UNIX-based platform installation. The `httpd.conf.sample` file was derived from the `httpd.conf.default` file with modifications for MPE. This is the file you should edit.

```
shell> cd /APACHE/PUB/conf
shell> cp mime.types.default mime.types
shell> cp magic.default magic
shell> cp httpd.conf.sample httpd.conf
shell> vi httpd.conf
```

1. Uncomment and replace "yourserver.com" with your own server's name for the `ServerAdmin` and `ServerName` directives. Review the Configuration Recommendation below. You may also choose to make other changes. The Apache Server Project web site, <http://www.apache.org/docs>, has information about all of the configuration directives.

All lines beginning with `#` are comments and are ignored by the HTTPD program file. Also note that the content of the configuration files is case sensitive.

Changes to the global configuration files do not take effect until the web server is started (or stopped and restarted if it is already running).

2. Verify the configuration file. It is a good idea to verify your configuration files before trying to start the web server. This verification is for syntax checking only.

```
shell>/APACHE/PUB/HTTPD -t
```

Configuration Recommendation

We strongly recommend setting `KeepAlive` to `off` in `http.conf`. With `KeepAlive` on, sockets may eventually be used up. When this happens, Apache will appear to be hung and will need to be restarted.

Edit the JHTTPD Job Stream File

The JHTTPD Job Stream File is used to run the HTTPD web server program in standalone mode with your local timezone:

```
!job jhttpd,www.apache,pub;outclass=,2
!setvar TZ `PST8PDT`
!run HTTPD;info='-f /APACHE/PUB/conf/httpd.conf'
!eoj
```

The `timezone` variable, `TZ`, should be set to your local timezone (for example, `EST5EDT` for Eastern Daylight Time, `PST8PDT` for Pacific Daylight Time, and `MST7MDT` for Mountain Daylight Time). For more information about setting the timezone value, enter “`man timezone`” in the POSIX shell to read the POSIX help file.

To preserve the access permissions (ACD) on the JHTTPD file, you can edit the file using the secure feature of EDITOR:

```
:EDITOR
/t JHTTPD
/m ...
/set secure
/k
JHTTPD,UNN
JHTTPD ALREADY EXISTS - RESPOND YES TO PURGE OLD AND KEEP NEW
PURGE OLD?yes
/e
```

Start the HTTPD Web Server Program

The HTTPD web server program is started using the JHTTPD job file. Apache can be started from either the CI or the POSIX shell:

```
:STREAM JHTTPD.PUB.APACHE
or
:xeq sh.hpbin.sys -L
shell>callci "stream jhttpd.pub.apache"
```

Verify that Apache is Running

There are a number of ways to verify if the Apache web server is running or, if it is not, to isolate how far the startup process progressed.

After streaming the JHTTPD file, use `:SHOWJOB` to view the running job:

```
JOBNUM  STATE  IPRI  JIN  JLIST      INTRODUCED  JOB NAME
#J16    EXEC           10S LP      TUE 10:27A  JHTTPD,WWW.APACHE
```

Another method is to check server status using either `ps` from the POSIX shell or using `:SHOWPROC` at the CI. The parent process is the HTTPD process with the `-f` option.

```
:SHOWPROC; job=jhttpd,www.apache
D202    0:00.079      WAIT   J16    73    :RUN HTTPD;info='-f
/APACHE/PUB/conf/httpd.conf'
D238    1:25.371      WAIT   J16    97    (HTTPD.PUB.APACHE) -f
/APACHE/PUB/conf/httpd.conf
D202    0:00.066      WAIT   J16    78    (HTTPD.PUB.APACHE)
D202    0:00.072      WAIT   J16    76    (HTTPD.PUB.APACHE)
D202    0:00.069      WAIT   J16    93    (HTTPD.PUB.APACHE)
D202    0:00.090      WAIT   J16   106    (HTTPD.PUB.APACHE)
D202    0:00.050      WAIT   J16    99    (HTTPD.PUB.APACHE)
D202    0:00.055      WAIT   J16   107    (HTTPD.PUB.APACHE)
```

```
shell> ps -ef | grep HTTPD
```

```
      UID      PID      PPID  C   STIME TTY      TIME COMMAND
WWW.APACHE  720972  655457  0   Dec 31 ldev10  0:00 HTTPD.PUB.APACHE
WWW.APACHE  589902  655457  0   Dec 31 ldev10  0:00 HTTPD.PUB.APACHE
WWW.APACHE  3014749 655457  0   Dec 31 ldev10  0:00 HTTPD.PUB.APACHE
WWW.APACHE  655457  5963849 0   Dec 31 ldev10  1:23 HTTPD.PUB.APACHE info=-f
WWW.APACHE  4587619 655457  0   Dec 31 ldev10  0:00 HTTPD.PUB.APACHE
WWW.APACHE  262250  655457  0   Dec 31 ldev10  0:00 HTTPD.PUB.APACHE
WWW.APACHE  3407979 655457  0   Dec 31 ldev10  0:00 HTTPD.PUB.APACHE
```

If the HTTPD program is running, make sure you can also access files in the directory tree. Some pages to try are:

- **Your server's Home Page**

```
http://yourserver.com
```

- **MGR.APACHE's Home Page**

```
shell> cd /APACHE/PUB/public_html
shell> cp index.htm.sample index.html
http://yourserver.com/~MGR.APACHE
```

- **The Apache on-line documentation manual**

```
http://yourserver.com/manual
```

If you are unsuccessful in starting the HTTPD program, you can get more information about the problem by:

1. Looking at the output of the JHTTPD spoolfile
2. Checking the messages in the `/APACHE/PUB/logs/error_log` file

3. Trying to telnet to Apache's port (port 80 if you have not modified the port directive set in the `httpd.conf.sample` file) If you cannot telnet to Apache's port, then your problem is not with the web server. You are not yet reaching your HP 3000 box. A successful telnet connection should look something like,

```
telnet yourserver.com 80
Trying...
Connected to yourserver.com.
Escape character is '^]'.

```

Stopping Apache

Apache is stopped by aborting the Apache job from either the CI or the POSIX shell. Use `SHOWJOB` to view the Apache job:

```

JOBNUM  STATE  IPRI  JIN  JLIST      INTRODUCED  JOB NAME
#J16    EXEC           10S LP      TUE 10:27A  JHTTPD,WWW.APACHE
:ABORTJOB #J16 or :ABORTJOB JHTTPD,WWW.APACHE

```

Apache can also be stopped using `:ABORTJOB` from the POSIX shell

```

:xeq sh.hpbin.sys -L
shell>callci "abortjob #j16" or shell>callci "abortjob jhttp,www.apache"

```

Logging

Apache error logging is useful when trying to start Apache as well as for monitoring a running web server. Apache logs errors into a log file called the `error_log` by default. This log file resides in the `/APACHE/PUB/logs` directory.

The number of messages logged in the error log is set by the `LogLevel` directive in the `httpd.conf` file. Possible values for `LogLevel` (by increasing significance) are `debug`, `info`, `notice`, `warn`, `error`, `crit`, `alert`, `emerg`. When a particular level is specified, messages from all other levels of higher significance are reported as well. For example, when `LogLevel` is "info", then messages with log levels of `notice`, `warn`, and up to `emerg` are also posted.

By default, the level is set to `warn`. Using a level of at least `crit` is recommended.

Version

The current version of Apache for MPE/iX is based on Apache 1.3.4 from the Apache Software Foundation. In addition to the `http` core that is the heart of the Apache code, 28 more Apache modules have been compiled into the web server program for a large feature set. A full description of each of these modules is included in the on-line manual document pages.

```

shell> /APACHE/PUB/HTTPD -v
Server version: Apache/1.3.4 (MPE-1.3.4-v1.0)
Server built:   Aug 24 1999 16:11:19
shell> /APACHE/PUB/HTTPD -l
Compiled-in modules:
    http_core.c

```

```
mod_env.c
mod_log_config.c
mod_log_agent.c
mod_log_referer.c
mod_mime_magic.c
mod_mime.c
mod_negotiation.c
mod_status.c
mod_info.c
mod_include.c
mod_autoindex.c
mod_dir.c
mod_cgi.c
mod_asis.c
mod_imap.c
mod_actions.c
mod_speling.c
mod_userdir.c
mod_alias.c
mod_access.c
mod_auth.c
mod_auth_anon.c
mod_cern_meta.c
mod_expires.c
mod_headers.c
mod_usertrack.c
mod_unique_id.c
mod_setenvif.c
```

Additional Documentation

Much of the public information available on Apache can be used for administrating Apache on MPE/iX. This especially applies to the description and usage of the over 128 different Apache configuration directives.

Sources for additional information include,

- The Apache on-line manual pages included as part of MPE FOS at <http://yourserver.com/manual/index.html>
- The Apache Software Foundation's on-line documentation at <http://www.apache.org/docs>
- Apache books, published by various publishers, such as O'Reilly and Associates, Inc. and IDG Books Worldwide, Inc.

For downloadable software to enhance your web site (perl, sendmail, python, etc.), visit MPE's external Jazz web server at <http://jazz.external.hp.com>

Introducing LDAP C-SDK/iX

by Joanne Gong
Commercial Systems Division

Lightweight Directory Access Protocol (LDAP) is an open standard which defines the communication protocol between an LDAP client and an LDAP server. LDAP C-SDK/iX is a set of client APIs in C language for MPE/iX applications to access X.500-based network directories.

LDAP C-SDK/iX enables internet directory accesses from HP 3000. An MPE client application can use an LDAP API to access LDAP directories on any LDAP server regardless of their platforms. Applications such as network management, corporate directory, yellow page/white page name search, mail and security services can all share directories across heterogeneous platforms.

An LDAP client connects to an LDAP server through SOCKET. Since LDAP is a platform-independent open standard, LDAP C-SDK/iX can connect to any standard-compliant LDAP server running on any platform. MPE/iX supports only the LDAP client, not the LDAP server. HP-UX does support the LDAP server.

LDAP C-SDK/iX is ported from the open source of Netscape Directory SDK Version 3.0 for C language. It is implemented with LDAPv3 standard. Although it can connect to the LDAPv2-based server, features and extensions introduced in LDAPv3 are not recognized by the LDAPv2-based server.

LDAP C-SDK/iX is automatically installed with the FOS in the `/usr/local/ldapsk` directory. It does not require additional installation from the customer. The home directory of LDAP C-SDK/iX, `/usr/local/ldapsdk/`, is comprised of:

- `lib/` containing API libraries `libldap.a` and `liblber.a`
- `include/` containing header files `ldap.h` and `lber.h`
- `tools/` containing command-line interface for simple query or testing
- `examples/` containing sample source code of LDAP client applications for developing LDAP-enabled applications

AUTOINST also creates the following symbolic links:

```
/usr/lib/libldap.a, /usr/lib/liblber.a, /usr/include/ldap.h  
/usr/include/lber.h
```

so that the API libraries and header files of LDAP C-SDK/iX can be found by the compiler automatically.

Detailed information on the features and API specifications also included in LDAP C-SDK/iX can be found in Netscape Directory *SDK 3.0 for C Programmer's Guide* at <http://developer.netscape.com/docs/manuals/dirsdk/csdk30/index.htm>.

MPE/iX Software Developer's Kit (SDK) for Java, Version 1.2.2

*by Mike Yawn,
Commercial Systems Division*

A major new release of Java, Version 1.2.2, is included with this release of MPE/iX. This article briefly describes some of the changes. One change is the name; the former Java Developer's Kit (JDK) is now called the Software Developer's Kit (SDK) for Java. You will also find in the literature that SDK 1.2 is frequently referred to as the Java 2 Platform. This is done merely to emphasize the large increase in functionality over JDK 1.1.x (which was presumably, although not explicitly, the Java 1 Platform). The Java 2 Platform does not indicate an SDK version of 2.0. Feel free to be confused by this.

This article does not attempt to cover all the new features of SDK 1.2, of which there are hundreds. Instead, it covers a small number of changes that are visible to the casual user. Programmers wishing to take advantage of all the new APIs and features can obtain a list at new features in the Java 2 platform at

<http://java.sun.com/products/jdk/1.2/docs/relnotes/features.html>.

Users just getting started with Java and wanting to learn the language may want to reference the tutorial at <http://java.sun.com/docs/books/tutorial/index.html>.

HTML format documentation was previously included as part of the JDK release. Because of the amount of disk space consumed by these files, they are now separately downloadable and installable from CSY's Jazz web server. As a result, you have the option of not installing the files at all, or of installing them on a different system if you do not run web server software on your HP 3000. The files can be downloaded from

<http://jazz.external.hp.com/src/java>. Navigate to the JDK 1.2 download page and select the HTML Documentation Files package. You can also browse the documentation on-line at <http://java.sun.com/products/jdk/1.2/docs/index.html>.

Previous versions of Java delivered for the HP 3000 included a Motif implementation of the Abstract Windowing Toolkit (AWT). The files comprising the AWT consumed a large amount of disk space, and feedback indicated that few users found the Motif implementation useful. We have discontinued distributing the AWT as part of our Java offering. For applications that require AWT functionality, the Remote AWT for Java technology from IBM can be used to allow graphical applications running on the HP 3000 to have their graphical displays redirected to a client system that also runs the Remote AWT software. The Remote AWT software is available at no charge from

<http://www.alphaworks.ibm.com/tech/remotewawtforjava>.

As with previous Java versions, you must stream `JINSTJDK.INSTALL.JAVA` to actually install the SDK on your system.

One minor but noticeable change in this release is in the directory structures for the Java product. The top level directory is still `/usr/local/java/<version>/`, with `<version>` being `jdk1.2.2` in this release. As in previous releases, there are subdirectories `/bin` and `/lib` at this level, containing binaries and libraries, respectively. New at this level is a `/jre` (for Java Runtime Environment) subdirectory, which is also subdivided into `/bin` and `/lib` subdirectories. This change is transparent to users, who should still use

`/usr/local/java/latest/bin` as the location for executable programs.

Further subdivision has happened in the libraries hierarchy; there is now a `/classic` subdirectory where the executable library (XL) containing most JVM functionality resides. This is in preparation for supporting multiple Java Virtual Machines running in the same directory hierarchy in a future release.

Certain command line options to Java that have been supported in previous releases are actually nonstandard, and not guaranteed to be available on all VM implementations. To help identify such nonstandard options, they now must be prefixed by `-X`. As an example, the `-verbose` option is standard, and thus is specified as `-verbose` on both 1.1.x and 1.2.x versions. `-nojit` is nonstandard, so `-nojit` works on the 1.1.x releases, but must be specified as `-Xnojit` on 1.2.x releases. Typing `java` with no options or filenames causes a list of standard options to be printed. Typing `java -X` causes a list of nonstandard options to be printed.

Releases prior to JDK 1.1.7 included a `JAVAUDC` UDC file in `PUB.JAVA`. In JDK 1.1.7, the UDC file was replaced by two command files in `PUB.SYS`: `JAVA` (to invoke the JVM) and `JAVAC` (to invoke the `javac` compiler).

If you have previously installed a web download version of **JAVA 2** on your HP 3000, the installation will have created command files named `JAVA2` and `JAVA2C` in `PUB.SYS`. With SDK 1.2.2, these files are no longer required and you can purge them. The `JAVA` and `JAVAC` command files distributed with this release will work correctly with both older JDK1.1.x versions of Java, and newer 1.2.x versions.

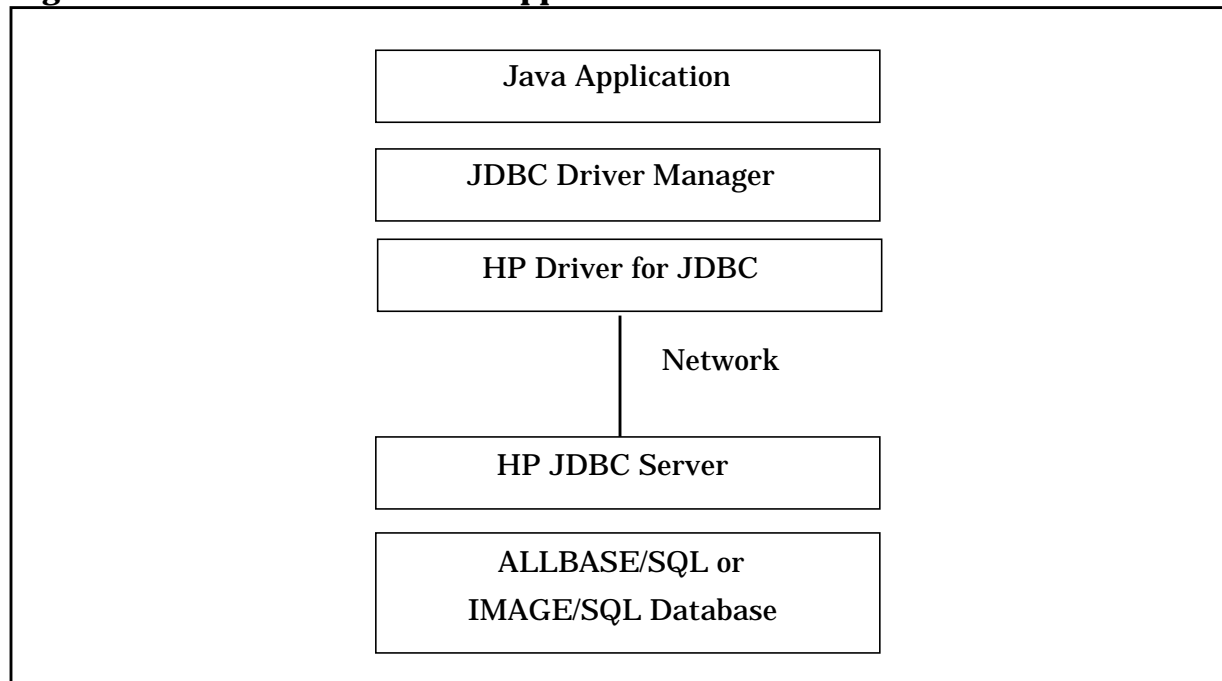
HP Driver for JDBC

Shu-Feng Wei
Commercial Systems Division

Java Database Connectivity (JDBC) is a Java API that enables development of Java applications and applets with a wide range of relational databases. It consists of a set of classes and interfaces written in Java. JDBC allows developers to write database applications using a standard, pure Java API.

A typical JDBC Java application consists of a Java application or applet, the JDBC Driver Manager, a vendor specific JDBC driver, and a database. The JDBC Driver Manager is provided with the Java SDK and its primary function is to load and register the vendor-specific JDBC driver with the Java applications and then get out of the way. The following diagram shows the various components of a client-server JDBC application.

Figure 5-1. JDBC Client-Server Application



The HP Driver for JDBC is a vendor specific JDBC Driver that allows Java applications to connect to ALLBASE/SQL or IMAGE/SQL databases on MPE/iX or ALLBASE/SQL on HP-UX. The HP Driver for JDBC is a Type 3 (Network-Protocol) driver, meaning the driver translates the JDBC API into a DBMS-independent protocol on the client-side, and then translates to the ALLBASE/SQL protocol on the server. The driver components on the client-side are written in 100% Pure Java, which provides the complete compatibility with all Java Virtual Machines on all platforms.

HP JDBC Components

There are three components supplied with the HP JDBC product, the JDBC Driver, the JDBC Monitor, and the JDBC Server. The user is responsible for writing a Java application or applet that uses JDBC on the client, as well as providing the ALLBASE/SQL or IMAGE/SQL database on the server.

- The **HP Driver for JDBC** is a set of Java classes that implement the `java.sql.*` interfaces and provide an implementation of a JDBC driver that can communicate with an ALLBASE/SQL or IMAGE/SQL database. The HP Driver for JDBC typically resides on the client side of the user application. It provides the translation from the Java language and the JDBC API to the HP proprietary network protocol.
- The **JDBC Monitor** is a component that is installed on the JDBC server host that manages all client JDBC Driver connections to the server host. It is typically started as a daemon when the server machine is booted. All JDBC client connections are made through the JDBC Monitor. The monitor performs validation of the userid and password that are passed in the client connection request and spawns JDBC Server processes to serve each of the client connections. Once the server process is spawned, the monitor returns to wait for the next client connection.
- The **JDBC Server** is the server process that is spawned by the JDBC Monitor to service a client connection. It handles the translation from the HP proprietary network protocol to the ALLBASE/SQL calls. There is at least one JDBC Server process for each client connection to the server host. More than one JDBC Server process may be used to handle multiple client statements using the same connection. This component also handles the translation from JDBC SQL to ALLBASE SQL and conversion of the database data from ALLBASE/SQL format to JDBC format.

Both the JDBC Monitor and the JDBC Server must be installed on the same host where the ALLBASE/SQL or IMAGE/SQL databases reside.

Requirements

Java Requirements

The HP JDBC Client components (the JDBC driver itself) require a Sun-compliant JDK version 1.1 and above, which includes a JDBC version 1.2. Install the JDK from Sun or from your platform vendor. Individual platform vendors may have their own requirements for the platform host. For example, Java only runs on HP-UX 10.20 and MPE/iX 6.0 and above, HP-UX 9.x and MPE/iX 5.5 is not supported.

Java and JDK components only need to be installed on the client platform. The JDBC server platforms use native components and must be installed on the same host where the ALLBASE/SQL or IMAGE/SQL databases reside.

HP-UX Server Requirements

The HP JDBC Server components require HP-UX version 10.20 or greater. HP-UX components are only required if your JDBC Server platform is the HP-UX operating system.

MPE/iX Server Requirements

The HP JDBC Server components require MPE/iX version 6.0 or greater. MPE/iX components are only required if your JDBC Server platform is the MPE/iX operating system.

ALLBASE/SQL or IMAGE/SQL Requirements

The HP JDBC Server components require an ALLBASE/SQL G3.01 or IMAGE/SQL G3.00 or greater.

Installation

HP-UX Server Components

The installation of the JDBC Server components must be done by a system administrator who has “root” capability on the server host where the database resides. A temporary directory, `/tmp/jdbc`, is used to stage the user manual, installation scripts, and product tar file.

The server components are distributed as a UNIX tape archive (TAR) file, `hpjdbc_XXX.tar`, where `XXX` represents the release number of the product.

Once the tar file is copied into a staging directory, `/tmp/jdbc`, extract the installation shell script from the tar file:

```
$ tar xf hpjdbc_XX.tar install.sh
```

MPE/iX Server Components

The installation of the HP JDBC product must be done by a system administrator on the server host where the database resides.

The components of HP JDBC are distributed as an archive file, `HFSFILES.JDBC.SYS`. Stream the JDBC installation job `I00IJDBC`, to perform the actual installation:

```
:stream I00IJDBC.JDBC.SYS
```

The installation script extracts the server files and client archive files and set up the JDBC monitor startup scripts.

Java Client Components

The HP JDBC product must first be installed on the server host before the client can be installed. This is because the client files are bundled with the server product.

The HP Driver for JDBC components consist of the Driver Java class files and a sample JDBC client source file. These are the only components required on the client-side. The driver components are packaged in three formats (the same files contents are in each package), UNIX tape archive (TAR), Java archive (JAR), and Windows ZIP (ZIP).

Setting up HP Driver for JDBC client involves:

- Downloading the appropriate HP JDBC archive file to a temporary directory.
- Extracting the HP JDBC archive file onto your client platform.

Extracting the HP Driver for JDBC Class Files The HP Driver for JDBC class files must be installed in your Java class path so that the Java compiler and the Java class loader can find them.

Example on the Win32 Platform

```
CLASSPATH=C:\JDK1.1.4\LIB;
```

You would then install the HP Driver for JDBC class files in the directory:

```
C:\JDK1.1.4\LIB\
```

Example on HP 3000

```
CLASSPATH=/usr/local/java/latest/lib:.
```

In order to use the HP Driver for JDBC client on the MPE/iX system, you need to have the JAVA/iX installed. JAVA/iX is shipped with MPE/iX starting in Release 6.0.

The following sub-directories for the Driver class files should be automatically created:

```
com/hp/jdbc/allbase
```

```
com/hp/jdbc/allbase/samples
```

Configuring HP JDBC Server

The HP JDBC Server components behave the same on both the HP-UX and MPE/iX platforms. Thus both the configuration file and log file are the same (except for filenames).

The following example shows the standard HP JDBC default server configuration file, `servcfg`. This is a text file that resides on the server host in the same directory as the HP JDBC Server executable files and can be used to alter some of the server's behavior.

```
LOGFILE /opt/allbase/jdbc/logs/servlog  
TIMEOUT 7200
```

```
#LOGGING ERROR  
#LOGGING CONNECTION  
#LOGGING INFO  
#LOGGING WARNING  
#LOGGING FATAL  
#LOGGING IN  
#LOGGING OUT  
#LOGGING INHEX  
#LOGGING OUTHEX  
#LOGGING DEBUG  
#LOGGING TIMESTAMP  
LOGGING NONE  
# If LOGGING NONE is not commented out, it must be the  
# last line for it to have the desired effect.
```

The first configuration option `LOGFILE` refers to the path and filename for the HP JDBC Server log file. On HP-UX this is normally set to `/opt/allbase/jdbc/logs/servlog`. On MPE/iX this is normally set to `SERVLOG.JDBC.SYS`.

The second configuration option `TIMEOUT` specifies the number of seconds that the HP JDBC Server can remain idle before terminating.

The next set of lines are the `LOGGING` levels used to determine the type of information to be logged to the server log file. The meanings of the various logging levels are discussed in the Troubleshooting Section, under “Server Logging” of the *HP Driver for JDBC User’s Manual*.

The server configuration file is read in each time a new server process is started. Thus once the file is edited and saved, the changes takes place beginning with the next server process.

If logging is turned on in this file, it is in effect for all subsequent server processes.

Starting and Stopping HP JDBC Monitor

HP JDBC Monitor is the server-side counterpart of HP Driver for JDBC, present on the client-side. For a JDBC application to connect to a database, the JDBC Monitor must be running on the database server. JDBC Monitor can be started manually or automatically at system startup by including the JDBC Monitor start commands in the system startup scripts.

HP-UX Monitor Startup and Shutdown The startup and shutdown of the HP-UX JDBC Monitor is normally controlled by the system startup and shutdown scripts. Thus as long as the machine is up and running, the JDBC Monitor will also be up.

The only time the Monitor should ever be brought down is to install a newer version. In this case, use the `monctrl` command to kill the Monitor process. Do not use the UNIX `kill` command, as this could render the server unable to start a new Monitor process. The `monctrl` command to shut down the monitor is:

```
monctrl kill [portnumber]
```

The `portnumber` parameter is necessary only if the monitor you wish to shut down is not running on the default port number 31700.

To restart the monitor after it has been accidentally terminated or was shut down, use the `monctrl` command to start it up:

```
monctrl start [portnumber]
```

Again, the `portnumber` parameter is necessary only if you wish to start the monitor on a port number other than the default port number 31700.

You cannot start a monitor on the same port number as a currently running monitor. You also can not restart a monitor on its original port number until all child processes that were spawned by the previous monitor process are terminated, and the port released. If the monitor is intentionally terminated or accidentally terminates, all of its child processes must be terminated before it can be restarted. For this reason, the monitor *must not* be terminated by using the UNIX `kill` command. Always use the `monctrl` tool to kill the monitor. The tool searches out all the child processes and kill them first, before killing the monitor.

MPE/iX Monitor Startup and Shutdown The startup and shutdown of the MPE/iX HP JDBC Monitor is normally done by the startup and shutdown stream jobs, `JSTRIMON` and `JSTOPMON`, which are normally included in MPE/iX system startup and shutdown scripts. Thus as long as the machine is up and running, the JDBC Monitor will also be up.

The only time the monitor should ever be brought down is to install a newer version. In this case, use the `JSTOPMON` stream job or `ABORTJOB` command to kill the monitor process. The `ABORTJOB` command to shut down the monitor is:

```
: ABORTJOB #JXX
```

where `XX` is the job number.

To restart the monitor after it was shut down or accidentally terminated, use the `JSTRIMON` stream job to start it up.

```
: STREAM JSTRIMON.JDBC.SYS
```

On the MPE/iX platform, each client connection is serviced by a separate process that is in the same session as the monitor. Thus, if the monitor is terminated or dies, the existing server process already serving the client requests also dies. This could result in client applications experiencing dropped connections.

Using the HP Driver for JDBC

Java Class Path

Before the HP Driver for JDBC can be used, the Java class path must be set to include the HP Driver for JDBC class files. Without this setting, both the Java compiler and the Java Run Time Environment is not able to locate the HP Driver for JDBC. Please refer to your *JDK* documentation for more explicit details as to how to set the Java class path. In general, the Java class path must include the directory in which the HP Driver for JDBC class files have been installed.

If the HP Driver for JDBC class files are installed in:

```
/opt/java/lib/com/hp/jdbc/allbase
```

then the Java class path must include the directory:

```
/opt/java/lib
```

In most cases, this would mean that the `CLASSPATH` environment variable must be set to something resembling:

Example on the HP 9000

```
CLASSPATH=/opt/java/lib:.
```

Sample JDBC Client

The HP Driver for JDBC comes with the source code to two JDBC applications, a sample client and a simple client. Both can be used to test the installation of JDBC components on the client and server.

The simple client is called `SimpleClient` and is a bare-bones text-based application that makes a connection to an `ALLBASE/SQL` or `IMAGE/SQL` database, and allows the user to send SQL statements and retrieve the results. There are no frills in the application, to make the code as simple as possible. This client should mainly be used for educational purposes. The rest of this discussion focuses on the more robust application.

The sample client is called `SampleClient` and is a more robust version of the simple client.

It shows how to use a dialog box to obtain user input, and also formats the result set output a lot better.

The first step to using the sample client is to build it by compiling the Java source code. Normally, you would do this by using the Java compiler command (the command you use may be different, depending on your JDK):

NOTE If you are using JDK1.2.1 or later, add the following line in `SampleClient.java` before the call to `class.forName`:

```
DriveManager.setLogStream(System.Out);
```

```
javac SampleClient.java
```

To run the sample client, follow the instructions provided with your Java SDK for your client platform. Make sure that your Java `CLASSPATH` includes the directory in which the HP Driver for JDBC class files have been installed, so that the Java Virtual Machine is able to load them. For most platforms, the command to run the sample client is:

```
java SampleClient [-w] [-t]
```

Follow the prompts given by the application to connect to your database and to execute SQL statements.

The sample client can be used as a starting point for developing your own applications or for troubleshooting connection problems with the database.

Loading the HP Driver for JDBC

The Java method that is used to load all JDBC drives is the `class.forName` method. To load the HP Driver for JDBC, the code is:

```
class.forName("com.hp.jdbc.allbase.JdbcDriver");
```

This loads the HP Driver for JDBC and registers it with the JDBC Driver Manager. Once a driver has been loaded and registered with the Driver Manager, it is ready to be used to connect to a database.

URL Syntax (Including User Name and Password)

The first connection method specifies all connection parameters, including the user name and password, in an URL string.

```
java.sql.DriverManager.getConnection(url)
```

where:

```
String url =  
    "jdbc:allbase://host[:port]/database?UID=uid&PWD=pwd"  
    "[&TRACE=trace]";
```

host Name or IP address of the sever host.

port Optional port number on which the JDBC Monitor is listening. If not specified, the default part number 31700 is used.

database ALLBASE/SQL or IMAGE/SQL database name.

<i>uid</i>	Server host userid that is authorized to access the database.
<i>pwd</i>	Server host password that matches the user id provided above.
<i>trace</i>	Optional trace values separated by the vertical bar “ ” character. For more information on tracing, see the Troubleshooting section.

NOTE This connection method may result in the user name and password being logged in various places, as URLs are commonly logged on various web servers and proxy servers. The URL is also logged if tracing is turned on. For this reason, the second connection method is preferred.

URL Syntax (Without User Name and Password)

The second connection method specifies the user id and password as method arguments, so this information is not present in the URL.

```
java.sql.DriverManager.getConnection(url, uid, pwd)
```

where:

```
String url =
    "jdbc:allbase://host[:port]/database[?TRACE=trace]";
```

```
String uid = "uid";
```

```
String pwd = "pwd";
```

<i>host</i>	Name or IP address of the sever host
<i>port</i>	Optional port number on which the JDBC Monitor is listening. If not specified, the default part number 31700 is used.
<i>database</i>	ALLBASE/SQL or IMAGE/SQL database name.
<i>uid</i>	Server host userid that is authorized to access the database.
<i>pwd</i>	Server host password that matches the user id provided above.
<i>trace</i>	Optional trace values separated by the vertical bar “ ” character. For more information on tracing, see the Troubleshooting section.

ALLBASE/SQL Specifics

ALLBASE/SQL to JDBC Data Type Mapping

Table 5-1. shows what the HP Driver for JDBC reports as the JDBC data type for each ALLBASE/SQL data type. These are the `java.sql.Types` values that are returned from the `java.sql.ResultSetMetaData.getColumnType` method. An “X” in the column

indicates the data type mapping.

Table 5-1. Data Type Mapping

JDBC DATA TYPES (vertical)	T I N Y I N T	S M A L L I N T	I N T E G E R	B I G I N T	R E A L	F L O A T	D O U B L E	D E C I M A L	N U M E R I C	B I T	C H A R	V A R C H A R	L O N G V A R C H A R	B I N A R Y	V A R B I N A R Y	L O N G V A R B I N A R Y	D A T E	T I M E	T I M E S T A M P
ALLBASE/SQL DATA TYPES (horizontal)																			
SMALLINT (16-bits)		X																	
INTEGER (32-bits)			X																
REAL					X														
FLOAT(1...24)					X														
FLOAT(25...53)							X												
DOUBLE PRECISION							X												
DECIMAL								X											
NUMERIC									X										
CHAR											X								
VARCHAR												X							
DATE																	X		
TIME																		X	
DATETIME																			X
INTERVAL											X								

ALLBASE/SQL to JDBC Data Type Conversions

Table 5-2. shows the supported data type conversions between ALLBASE/SQL and JDBC. For conversions from JDBC to Java, please refer to a JDBC book or the Java JDBC documentation. Those conversions are generic to all JDBC Drivers.

Suggested conversions are denoted by a capital letter “X” in the conversion grid. Conversions which are supported, but which may result in a loss of precision, overflow, or

rounding, are denoted by a lowercase letter “x” in the conversion grid.

Table 5-2. Data Type Conversions

JDBC DATA TYPES (vertical)	T I N Y I N T	S M A L L I N T	I N T E G E R	B I G I N T	R E A L	F L O A T	D O U B L E	D E C I M A L	N U M E R I C	B I T	C H A R	V A R C H A R	L O N G V A R C H A R	B I N A R Y	V A R B I N A R Y	L O N G V A R B I N A R Y	D A T E	T I M E	T I M E S T A M P
ALLBASE/SQL DATA TYPES (horizontal)																			
SMALLINT (16-bits)	x	X	X	X	X	X	X	X	X	x	X	X	X	x	x	x			
INTEGER (32-bits)	x	x	X	X	x	x	x	X	X	x	X	X	X	x	x	x			
REAL	x	x	x	x	X	X	X	X	X	X	X	X	X	x	x	x			
FLOAT(1...24)	x	x	x	x	x	x	X	X	X	X	X	X	X	x	x	x			
FLOAT(25...53)	x	x	x	x	x	x	X	X	X	X	x	x	x	x	x	x			
DOUBLE PRECISION	x	x	x	x	x	x	X	X	X	x	X	X	X	x	x	x			
DECIMAL	x	x	x	x	x	x	x	X	X	x	X	X	X	x	x	x			
NUMERIC	x	x	x	x	x	x	x	X	X	x	X	X	X	x	x	x			
CHAR	x	x	x	x	x	x	x	x	x	x	X	X	X	x	x	x			
VARCHAR	x	x	x	x	x	x	x	x	x	x	X	X	X	x	x	x			
DATE											X	X	X	x	x	x	X		x
TIME											X	X	X	x	x	x		X	x
DATETIME											X	X	X	x	x	x	x	x	X
INTERVAL											X	X	X	x	x	x			

The conversion of any ALLBASE/SQL data type to the `java.sql.Types.BIT` data type is such that only the value of zero, in either numeric or character format, is converted to the `java.sql.Types.BIT` value of 0. All other values are converted to the `java.sql.Types.BIT` value of 1. Thus only the integer value 0, the floating-point value 0.0, the decimal value 0 (not 0.0), and the character string “0” is converted to a bit value of 0. Everything else is converted to a bit value of 1.

The conversion of an ALLBASE/SQL data type to a JDBC data type that has a smaller degree of precision (such as conversion from ALLBASE/SQL `INTEGER` to `java.sql.Types.SMALLINT`) follows the Java VM rules of casting one data type to another. This may result in different values on different Java platforms. There is also no warning generated when this occurs. It is up to the application developer to choose the appropriate JDBC data type.

The conversion of an ALLBASE/SQL character data type to a JDBC numeric data type uses the Java numeric conversion routines and any necessary numeric casting. Thus this

could result in a `java.lang.NumberFormatException` being shown when the conversion is performed. Again, the application developer should take the necessary precautions.

Unsupported ALLBASE/SQL Data Types

The ALLBASE LONG BINARY, and LONG VARBINARY data types are currently not supported by the HP Driver for JDBC.

Acceptable SQL Syntax

The HP JDBC Server accepts either ODBC SQL statement syntax or ALLBASE/SQL statement syntax. All SQL statements are first parsed for ODBC 2.0 SQL syntax. If the statement conforms to the ODBC 2.0 SQL syntax, it is translated to ALLBASE/SQL before being passed to the ALLBASE/SQL DBMS. If the statement does not conform to the ODBC 2.0 SQL syntax rules, it is assumed to be an ALLBASE/SQL statement, and is passed without modification to the ALLBASE/SQL SQL DBMS.

Unsupported ALLBASE/SQL SQL Statements

Table 5-3. shows a list of unsupported ALLBASE/SQL statement types in HP JDBC. Note that in many cases, JDBC provides a standardized method of performing the same action. For example, the ALLBASE/SQL “COMMIT WORK” statement is not supported, since the user should be using the `java.sql.Connection.commit` method.

Table 5-3. Unsupported Statements

ADVANCE	BEGIN DECLARE SECTION	BEGIN WORK
CLOSE	COMMIT WORK	CONNECT
DECLARE CURSOR	DELETE WHERE CURRENT	DESCRIBE
DISCONNECT	END DECLARE SECTION	EXECUTE
EXTRACT	FETCH	INCLUDE
OPEN	PREPARE	RELEASE
ROLLBACK WORK	SET CONNECTION	SET SESSION
SET TRANSACTION	SETOPT	START DBE
STOP DBE	SQL EXPLAIN	UPDATE WHERE CURRENT
TERMINATE USER	WHENEVER	

Dynamic/Parameterized SQL Statements

The HP Driver for JDBC supports the use of parameterized SQL statements through the `java.sql.PreparedStatement` interface. The SQL statements must use a question mark (?) as the marker character for passing the parameters.

For example, a SELECT statement with parameters in the where clause would look like:

```
SELECT NAME, ADDRESS FROM ADDRBOOK WHERE NAME=?
```

Another example is an INSERT statement:

```
INSERT INTO ADDRBOOK(NAME, ADDRESS) VALUES (?,?)
```

Stored Procedures

The HP Driver for JDBC supports the following types of ALLBASE/SQL stored procedures through the `java.sql.CallableStatement` interface:

1. Procedures that return one or more result sets.
2. Procedures that take one or more input arguments.
3. Procedures that return one or more output arguments.
4. Procedures that have a return status value.

Troubleshooting

Client Tracing

Tracing of the HP Driver for JDBC client class files is invoked by adding tracing commands to the connection URL. The connection URL can be altered to both invoke tracing and to control the type of information that is traced. The tracing information is sent to a Java stream, which must be specified by the application using the `java.sql.DriverManager.setLogStream` method.

Both the tracing level and tracing output must be specified before any tracing can be done.

Note that the use of tracing impacts performance of the application. The greater the detail in the tracing, the slower the performance of the application.

The URL syntax to invoke tracing is:

```
"jdbc:allbase://server[:port]/database?TRACE=trace"
```

where *trace* is any vertical bar (|) separated combination of the values:

ARGUMENTS	Trace HP Driver for JDBC class method arguments. Input arguments and return values are all traced. Only the methods called by the application are traced.
ARGUMENTS_ALL	Trace all HP Driver for JDBC class methods that are called by both the application and the driver itself.
TIME	Include the time in HHMMSSFFF format on all tracing output lines where HH is the hour from 00 to 23, MM is the minute from 00 to 59, SS is the second from 00 to 59, and FFF is the millisecond from 000 to 999. The information appears in the third column of the tracing output.
DATE	Include the date in YYYYMMDD format on all tracing output lines where YYYY is the year from 0000 to 9999, MM is the month from 01 to 12, DD is the day of the month from 00 to 31. This information appears in the second column of the tracing output.
TIMESTAMP	Include both the date and the time on all tracing output lines. This is just a combination of the DATE and TIME tracing values.

OBJECT	Include the object hash code in the format HHHHHHHH on all tracing output lines where HHHHHHHH is the eight digit hexadecimal value of the hash code. This information appears in the first column of the tracing output.
THREAD	Include the executing thread name on all tracing output lines. This information appears as the fourth column of the tracing output.
DEBUG	Trace debugging statements from the driver. This produces a large amount of output, so it should be used sparingly.
CONNECTION	Trace driver connection information.
NETWORK	Trace driver network information.
ERROR	Trace driver errors.
WARNING	Trace driver warnings.
SQL	Trace SQL statements that are passed to the driver.
ALL	Trace everything.

For example, the URL to turn on tracing of all connections with time stamp information would be:

```
"jdbc:allbase://server/database?TRACE=CONNECTION|TIMESTAMP"
```

Do not use any space characters between tracing levels and the vertical bar separator as some browsers and web servers do not allow space characters in a URL.

NOTE Remember that the application must also set the tracing stream using the `java.sql.DriverManager.setLogStream` method.

Server Logging

The server configuration file controls logging of the JDBC Server. The following is a list of the various logging levels that can be set in the server configuration file.

NONE	Turns off all preceding logging levels.
CONNECTION	Turns on logging of connection information such as the time of the connection, the client IP address, and the name of the database.
IN	Turns on logging of input information coming from the client.
INHEX	Turns on logging of the hexadecimal dump of the input information coming from the client.
OUT	Turns on logging of output information being sent back to the client.
OUTHEX	Turns on logging of the hexadecimal dump of the output information being sent back to the client.
INFO	Turns on logging of miscellaneous information about the processing of the server. This has the potential to log a fair amount of data, which may

	affect performance.
TIMESTAMP	This causes each log entry to be prefaced by a timestamp of the form HH:MM:SS.mmm.
WARNING	Turns on logging of warning messages generated by the Server.
ERROR	Turns on logging of error messages generated by the Server.
FATAL	This is the same as specifying INFO, WARNING, and ERROR.
DEBUG	Turns on logging of a lot of debugging messages. This should only be done at the request of support personnel. This generates a large amount of logging information and can severely affect performance.

6 Announcing a Post 6.5 Release Patch

This chapter provides the following article about a software patch (B.G\$.01) which is being released separate from and following the 6.5 base release. The software patch is not part of the 6.5 base release, but will be available shortly after the 6.5 release.

This article provides information regarding mandatory changes which you must make to your system if you decide to install this patch. Please read this article carefully so that you understand the impact of this required migration of the `ACTINFO` file.

NOTE Before migrating your `ATCINFO` file of your DBE, make sure you backup your existing `ATCINFO` file and the matching DBE.

- IMAGE/SQL Date Mapping Enhancement

IMAGE/SQL Date Mapping Enhancement

CAUTION The DATE MAPPING enhancement for IMAGE/SQL will be available in the first IMAGE/SQL patch released on MPE/iX version 6.5. This enhancement requires migration of the ATCINFO file since the ATCINFO file format has been changed to keep track of additional date related information.

The IMAGE/SQL version number for 6.5 release is B.G3.01. Since this version does not include the DATE MAPPING enhancement, no migration is needed. However, once you apply the IMAGE/SQL patch with the version number of B.G4.01, or any version following the B.G4.01 version, you need to migrate all your ATCINFO files in order to use IMAGE/SQL.

Currently, users are experiencing some problems accessing various dates in TurboIMAGE databases through IMAGE/SQL. Some typical problems are:

- In TurboIMAGE, dates may be held in numeric or ascii fields. If no data exists in a data field, it will hold binary zeroes, and will be interpreted by SQL as zero rather than null. Any attempt to CAST this data into a DATE format will fail with an error, and the SELECT can not continue. Similarly, a failure in the validation of dates in the 3GL code will result in the same error, and the SELECT statement fails.
- Both FA3000 and MM3000 will be handling year 2000 and beyond by incrementing the year to 'A0' for the year 2000, 'B0' for year 2010 etc. This allows their applications to be able to distinguish the year 2000 and beyond with a 'YYMMDD' format. The users of FA3000 and MM3000 need a way for IMAGE/SQL to convert the above format year to SQL DATE and vice versa. To alleviate the problems mentioned above, IMAGE/SQL has been enhanced to allow some numeric or ascii fields to be mapped to DATE columns, and the users can specify what DATE format the data is in. The date formats supported are the same as the formats supported by MPE. Following is a description of the enhanced SPLIT and UPDATE TYPE commands in the IMAGESQL program.

IMAGE/SQL SPLIT command and UPDATE TYPE command have been enhanced to allow SIGNED or UNSIGNED mapped type if the mapped type is DECIMAL. If UNSIGNED keyword is specified, then all positive values will be unsigned. The default is SIGNED. If the mapped type is DATE, and the source type is P8, Z6 or Z8, SIGNED or UNSIGNED keyword can be used to indicate whether all positive values will be SIGNED or UNSIGNED. The default is SIGNED.

IMAGE/SQL UPDATE TYPE and SPLIT commands are also enhanced to allow users to specify FORMAT=BB1 or BB2. (BB1 is for Business Basic short decimal, and BB2 for Business Basic decimal). The NewMappedType must be FLOAT. An error will be returned if a value is not in the right range.

This enhancement also has a dependency on the Business basic patch BBRDXU7 for the new Business Basic conversion routines.

SP(LIT)

Divides a large mapped column into two or more smaller columns.

Syntax

```
SP[LIT] MappedTable.MappedColumn INTO
  NewMappedColumn:SourceType
      [FORMAT=FormatType[,lowvalue[,highvalue]]]
      [:MappedType [SIGNED ]
              [UNSIGNED]] [,...]
```

If MappedType is 'DATE', then FORMAT must be specified.

Example

```
(i) SPLIT TABLE1.COLUMN5 INTO NEWCOLUMN1:I4:CHAR(8),&
    NEWCOLUMN2:X20,&
    NEWCOLUMN3:K3:DECIMAL(15,0)
(ii) SPLIT TABLE2.COLUMN3 INTO NEWCOL1:I4:CHAR(8),&
    NEWCOL2:X20,&
    NEWCOL3:Z4:DECIMAL(4,0) UNSIGNED
(iii) SPLIT TABLE2.MYDATE INTO &
    MYDATE1:X6 FORMAT=DT35,'000000','999999':DATE, &
    MYDATE2:X6 FORMAT=DT36,'*****','&&&&&&':DATE
(iv) SPLIT TABLE2.DDATE INTO &
    DDATE1:X6 FORMAT=DT26,' 0',' 99':DATE,&
    DDATE2:X6 FORMAT=DT25,' ','-----':DATE
(vi) SPLIT TABLE2.MYDATE2 INTO &
    MYDATE2_1:I2 FORMAT=DT2,0,-1:DATE, &
    MYDATE2_2:I2 FORMAT=DT2,0,-1:DATE
(vii) SPLIT TABLE1.COLUMN5 INTO NEWCOLUMN1:I4:CHAR(8),&
    NEWCOUNML2:X20,&
    NEWCOLUMN3:K4:FORMAT=BB2:FLOAT
(viii) SPLIT TABLE2.MYBBDECS INTO &
    MYBBDEC_1:K2 FORMAT=BB1:FLOAT, &
    MYBBDEC_1:K4 FORMAT=BB2:FLOAT
```

U[PDATE] TYPE

Updates data type mapping information for a specified TurboIMAGE/iX data type or a specified Mapped Column.

Syntax

```
U[PDATE] TYPE {SourceType IN {*
                                {MappedTable}
                                [FORMAT=FormatType[,lowvalue[,highvalue]]]}
               {IN MappedTable.Col
               [FORMAT=FormatType[,lowvalue[,highvalue]]]}
               [TO NewMappedType [SIGNED ]
               [UNSIGNED]]
```

If NewMappedType is 'DATE', then FORMAT must be specified.

Examples

```
(i) UPDATE TYPE I4 IN COMPOSER
(ii) UPDATE TYPE IN COMPOSER.BIRTHDATE TO CHAR(18)
(iii) UPDATE TYPE IN TABLE1.UNSIGNEDZ4 TO DECIMAL(4,0) UNSIGNED
(iv) UPDATE TYPE IN TABLE2.MYDATE FORMAT=DT35,'000000','999999' TO DATE
(v) UPDATE TYPE X6 IN TABLE2 FORMAT=DT26,' 0',' 99' TO DATE
(vi) UPDATE TYPE IN TABLE2.MYDATE2 FORMAT=DT2,0,-1 TO DATE
(vii) UPDATE TYPE IN TABLE2.MYP8 FORMAT=DT15,0,-1 TO DATE UNSIGNED
(viii) UPDATE TYPE IN TABLE3.MYZ6 FORMAT=DT26,'*****','999999' &
        TO DATE SIGNED
(ix) UPDATE TYPE IN TABLE4.MYZ8 FORMAT=DT38,'-----','#####' & TO DATE
UNUNSIGNED
(x) UPDATE TYPE K2 IN TABLE2 FORMAT=BB1 TO FLOAT
(xi) UPDATE TYPE IN TABLE2.MYBBDEC FORMAT=BB2 TO FLOAT(53)
```

Format(TI)Type StorageType #Bytes Explanation

Format (TI)	Storage Type	Bytes	Explanation
DT1 (I4,J4,K4)	longint	8	MPE time-stamp (microseconds from 1970-01-01)
DT2 (I2,J2,K2)	integer	4	Upper 2 bytes: year next byte: month of year
DT3 (I2,J2,K2)	integer	4	Upper 2 bytes: year bottom 2 bytes: day of year
DT4 (I2,J2,K2)	integer	4	Upper 23 bits: # years from 1900 bottom 9 bits: day of the year. (analogous to the existing CALENDAR format.)
DT10 (I2,J2,K2)	integer	4	Seconds from 1970-01-01 (POSIX.1 time() format; valid through 2038-01-18)
DT14 (I1,J1,K1)	shortint	2	Upper 7 bits: #years from 1900 Lower 9 bits: day of the year (CALENDAR format; valid up to 2027-12-31)
DT15 (I2,J2,K2,P8)	integer	4	YYMMDD date
DT16 (I2,J2,K2,P8)	integer	4	MMDDYY date
DT17 (I2,J2,K2,P8)	integer	4	DDMMYY date
DT18 (I2,J2,K2)	integer	4	YYYYMMDD date
DT25 (X6,U6,K3,Z6)	ASCII	6	YYMMDD date
DT26 (X6,U6,K3,Z6)	ASCII	6	MMDDYY date
DT27 (X6,U6,K3,Z6)	ASCII	6	DDMMYY date
DT35 (X6,U6,K3)	ASCII	6	YYMMDD date YY:MM3000 date
DT36 (X6,U6,K3)	ASCII	6	MMDDYY date YY:MM3000 date
DT37 (X6,U6,K3)	ASCII	6	DDMMYY date YY:MM3000 date
DT38 (X8,U8,K4,Z8)	ASCII	8	YYYYMMDD date
BB1 (K2)		4	Business Basic SHORT DECIMAL type
BB2 (K4)		8	Business Basic DECIMAL type

Notes

- MM3000 dates are represented as in the MM3000 product which uses the ASCII "A" through "Z" for decades starting with the year 2000.
- lowvalue will be displayed as '9900-01-01',
highvalue will be displayed as '9999-12-31', and

invalid value will be displayed as '9901-12-31'

There will be no default values for lowvalue and highvalue. If lowvalue and/or highvalue are not provided, then the code checking for lowvalue/highvalue will not be exercised.

3. For integer DT types (DT1 - DT18), the lowvalue and the highvalue must be set to integer values.
4. For ASCII DT types (DT25 - DT38), the lowvalue and the highvalue must be set to ascii values (single quoted), 0 (for COBOL LOWVALUES), or -1 (for COBOL HIGHVALUES). Single quote character cannot be included in the lowvalue or highvalue ascii string. No padding will be done for the user. The user must provide ascii values with the correct length. For example, for DT25, the lowvalue and highvalue provided must be 6 characters long.
5. If the TURBO data type is P8, Z6 or Z8, and it is being mapped to DATE, then [UNSIGNED/SIGNED] keyword can be used to indicate whether positive data should be SIGNED or UNSIGNED when inserting data via IMAGE/SQL. The default is SIGNED.

Examples

1. Use IMAGESQL:

```
UPDATE TYPE IN mytable.mydate FORMAT=DT26,'*****','$$$$$$' TO DATE
```

2. Use ISQL

```
INSERT INTO mytable VALUES ('9900-01-01');  
This statement will PUT '*****' into the TURBO field mydate.  
  
INSERT INTO mytable VALUES ('9999-12-31');  
This statement will PUT '$$$$$$' into the TURBO field mydate.  
  
INSERT INTO date2.d2 VALUES ('9901-12-31','9901-12-31','bad date');  
Invalid DATE for TurboIMAGE: set 1, item 0, column 2, record 0.  
(DBERR 13512)  
Number of rows processed is 0  
  
SELECT mydate FROM mytable where mydate is null;  
returns "Number of rows selected is 0"  
i.e. nothing qualifies.  
  
SELECT mydate FROM mytable WHERE mydate = '9900-01-01';  
returns all records with TURBO value = '*****'  
(if index exist on the item, mydate, then CHAIN READ is used)  
  
SELECT mydate FROM mytable WHERE mydate = '9999-12-31';  
returns all records with TURBO value = '$$$$$$'  
(if index exist on the item, mydate, then CHAIN READ is used)  
  
SELECT mydate FROM mytable WHERE mydate = '9901-12-31';  
returns all records with TURBO value is not a valid date.  
(SERIAL READ is always used in this case)
```

3. How do you migrate your ATCINFO file?

NOTE Before migrating your ATCINFO file of your DBE, make sure you backup your existing ATCINFO file and the matching DBE.

```
: run imagesql.pub.sys  
> dbe yourdbe  
> migrate
```

4. How do you know that your ATCINFO file needs migration?

```
(1) : imagesql  
    >> dbe mydbe  
    >> base yourdb  
    >> display turbodb  
    Incompatible ATCINFO file version number, migration is needed  
    (ATCERR 2101).  
(2) isql=> connect to 'mydbe';  
    ATCINFO file contains incompatible version number, migration is  
    needed. (DBERR 13516)
```

5. How about going back to an older version?

Possible scenario: You restored the migrated database onto a system that is running an older version of IMAGE/SQL.

Symptoms:

```
(1) :imagesql  
    >>dbe mydbe  
    >>base yourdb  
    >>display turbodb  
    ATCINFO procedure error (ATCERR 32421,ATCSTAT 2101,  
    FSERR 0).  
(2) isql=> connect to 'mydbe';  
    IMAGE/SQL internal error 141, 2101, 0, 0. (DBERR 13554)
```

What to do:

Restore the ATCINFO file and the matching DBE from a backup before the migration.

Announcing a Post 6.5 Release Patch
IMAGE/SQL Date Mapping Enhancement

7 Product List—How to Order

You can order subsystem products through your HP Sales Representatives.

- To order products in the U.S., call HP Computer Organization Order Fulfillment Center at 1-800-386-1117 option 3.
- To order learning products manuals in the U.S., call HP Parts Director Ordering at 1-800-227-8164.
- To order products and/or manuals in countries outside the U.S., call your local Trade Parts Sales Organization.

You must specify the product number and release base (for this release, Release 6.5) when ordering. Installing the PowerPatch tape updates your system with the FOS enhancements, as well as subsystems that have patches.

You may order only the PowerPatch tape. However, installing the PowerPatch ONLY, updates your system only with the FOS enhancements—the subsystem products **must** be ordered and installed via the SUBSYS tape from MPE/iX Release 6.5. Customers with Response Center Support can contact the Response Center to order the PowerPatch tape.

NOTE MPE/iX Release 6.5 can be ordered only by mailing or faxing the order form sent to those customers with an appropriate support contract. It cannot be ordered by phone.

8 Product Release History

This chapter contains tables that provide information on the currently supported Commercial Systems MPE/iX releases and products, and the systems supported for the 6.5 Release..

Product Changes by Releases

The following table provides information on the currently supported Commercial Systems MPE/iX releases and products. Included are the MPE/iX release or SUBSYS VUF and a list of products introduced. It also provides information on significant changes made to a release. Enhancements, NPCONFIG Variable on NW Spooler, DLT4000/DLT7000 Differential Tape Drives, IMAGE/SQL Enhancement: P and Z Data Types, Java Database Connectivity (JDBC) Driver, Java Developer's Kit, Legato NetWorker Storage Node for MPE/iX, FTP Enhancements, and HP 3000 997 Large Memory Subsystem

Table 8-1. MPE/iX Product Releases

Release	SUBSYS	Date Code	Product(s) Introduced/Added
C.55.00		R3628	MPE/iX Release 5.5 (Non-Platform Release) HP Loader Dependent Libraries Subsystem Dump Facility HP Stage/iX HP Patch/iX TurboSTORE/iX 7x24 True-Online Backup HP Optical Disk Libraries: C1150B - 40GB C1160B - 80GB C1170B - 100GB TCP/IP Network Printer Support Telnet/iX Server DTS/TIO Dynamic Configuration Online System Device Configuration
5.5 PP 1	C.55.01	R3638	PowerPatch 1 based on Release 5.5 ALLBASE/SQL CAST function HP Telnet/iX—full functionality CCSY Access Server 979/x00KS support High Availability Disk Array support

Table 8-1. MPE/iX Product Releases

Release	SUBSYS	Date Code	Product(s) Introduced/Added
5.5 Exp. 2	C.55.02	R3715	Express 2 based on Release 5.5 CI enhancements VPlus enhancements DTS/TIO new functionality
5.5 Exp. 3	C.55.03	R3732	Express 3 based on Release 5.5 COBOL II/iX enhancements 100VG-AnyLAN Network Adapter introduced 100Base-T Network Adapter introduced ALLBASE/SQL release G2 ODBCLink/SE ODBC driver introduced TurboIMAGE/XL B-Tree and scalability enhancements IMAGE/SQL B-Tree enhancement 9 GB Disk Drive NMMGR Node Maintenance Manager enhancements New SCSI Disks: ST32272N, ST32272WC (2 GB) ST34572N, ST34572WC (4 GB) XP34361W (4 GB)
5.5 Exp. 4	C55.04	R3748	Express 4 based on Release 5.5 Year 2000 enhancements MPE/iX new date intrinsics 997/100-500 support Enhancements for: Transact V and Transact/iX Dictionary/3000 COBOL II/iX Inform/V HP ALLBASE/BRW QUERY/iX
5.5 PP 5	C55.05	R3813	PowerPatch 5 based on Release 5.5 989/100, 989/200, 989/400, 989/600 support 997/600, 997/800 support A3699A - HP-IB I/O Bus Converter New SCSI Disks: A4909A - 18GB, WD A5239A - 18GB, WD Enhancements for Image/SQL and TurboIMAGE/XL: Dynamic data set expansion for master sets Performance enhancement for TPI Image Row-level locking 80 GB Image data sets ODBCLink/SE E.56.12

Table 8-1. MPE/iX Product Releases

Release	SUBSYS	Date Code	Product(s) Introduced/Added
5.5 PP6	C.55.06		PowerPatch 6 based on Release 5.5 Bug Fixes only
5.5 Exp. 7	C.55.07	R3910	Express 7 based on Release 5.5 989/150, 989/250, 989/450, 989/650 support ALLBASE/SQL Release G3 New String Functions ALLBASE/BRW JCWs for year 2000 DLT4000/DLT7000 Differential Tape Support IMAGE/SQL Enhancement Entering data with P and Z data types SPLIT and UPDATE commands allow mapped type Legato NetWorker Storage Node NPCONFIG variable for NW Spooler
C.60.00	C.60.00	R3812	MPE/iX Release 6.0 (Platform Release) User-defined job Queue FTP enhancements Java for MPE/iX Samba/iX Performance enhancements System limit enhancements DNS BIND/iX
C.60.01	C.60.01	R3926	PowerPatch 1 based on Release 6.0 CI Enhancements NPCONFIG Variable on NW Spooler PATCH/iX Enhancements DLT4000/DLT7000 Differential Tape Drives IMAGE/SQL Enhancement: P and Z Data Types HP Driver FOR JDBC Java Developer's Kit Version 1.1.7B Legato NetWorker Storage Node for MPE/iX HP 3000 997 Large Memory Subsystem Support for 36 Gigabyte Disks
C.65.00	C.65.00		MPE/iX Release 6.5 (Platform Release) Support for large files Increased TCP connections Support for 511 Disks Support Tools Manager (STM) Enterprise Management Solution HP Secure Web Console Apache for MPE/iX LDAP C-SDK/iX NEWCI Command

Supported Releases

The naming conventions for the different types of releases have been changed slightly to clarify the type of release being discussed. The terms used to describe or refer to the releases are:

Mainline Release A mainline release involves the recompilization and reintegration of all software release components (FOS, SLT, and SUBSYS tapes). The release number is changed (for example, 6.0 or 6.5) and the update "UU" field of the V.UU.FF is changed. There are two types of mainline releases: *Platform* and *Non-Platform*.

Platform Release A platform release (previously also known as a "core" release) is a subset of a mainline release. Typically, the release number ends with a "0" such as 6.0. Platform releases are *automatically distributed* to all customers with support contracts.

Non-Platform Release A non-platform release (previously known as a "major" release) is a subset of the mainline release. The release number typically ends with a "5" such as 6.5. Non-platform releases must be *explicitly ordered* by customers.

Both platform and non-platform releases can be referred to as *mainline releases* when not discussing distribution or extended support life.

The following matrix provides information on the supported Commercial Systems MPE/iX mainline releases. It lists the currently supported releases and the SPUs they are supported on. The matrix also provides all known factory support termination dates. When a mainline release becomes unsupported, the factory will not provide any support services for that release. Online calls are not accepted and patches are not created; customers are advised to roll to a supported release.

- In general, HP will support the two most recent mainline releases.
- In order to facilitate customers moving to new releases, we will support the oldest release for at least six months after the new release ships. In other words, Release "N-2" will still be supported for at least six months after Release "N" ships.
- The release of an Express Release or a Powerpatch on any particular release does not extend its support life.

Given the rate with which we ship releases, this new strategy implies that any particular release will be supported for at least three years, and in most cases, most releases will be supported for a significantly longer period of time.

This new strategy is effective now. This implies that Release 5.5 will become obsolete as early as six months after Release 6.5 ships.

Table 8-2. Supported System Release Matrix

Supported Releases	Supported Systems	Support Termination Date
Release 5.5 (C.55.xx)	920*, 922*, 932*, 948*, 955**, 958*, 960**, 980/100, 980/200, 980/300, 980/400, 9x7, 9x7LX, 9x7RX, 9x7SX, 9x8LX, 9x8RX, 939KS, 939KS/020, 959KS/100, 959KS/200, 959KS/300, 959KS/400, 969KS/100, 969KS/200, 969KS/300, 969KS/400, 969KS/120, 969KS/220, 969KS/320, 969KS/420, 979KS/100, 979KS/200, 979KS/300, 979KS/400, 989/100, 989/200, 989/400, 989/600, 989/150, 989/250, 989/300, 989/350, 989/450, 989/500, 989/550, 989/650, 990CX 992/100CX, 992/200CX, 992/300CX, 992/400CX, 990DX, 992/100DX, 992/200DX, 992/300DX, 992/400DX, 991CX, 995/100CX, 995/200CX, 995/300CX, 995/400CX, 995/500CX, 995/600CX, 995/700CX, 995/800CX, 991DX, 995/100DX, 995/200DX, 995/300DX, 995/400DX, 995/500DX, 995/600DX, 995/700DX, 995/800DX, 996/80, 996/100, 996/200, 996/300, 996/400, 996/500, 996/600, 996/700, 996/800, 996/900, 996/1000, 996/1200, 997/100, 997/200, 997/300, 997/400, 997/500, 997/600, 997/800	6 months following next Mainline release after 6.0 (about end of year 2000)
* Support life ends as of 1/00 ** Support life ends as of 8/00		

Table 8-2. Supported System Release Matrix

Supported Releases	Supported Systems	Support Termination Date
Release 6.0 (60.xx)	920*, 922*, 932*, 948*, 955**, 958*, 960**, 980/100, 980/200, 980/300, 980/400, 9x7, 9x7LX, 9x7RX, 9x7SX, 9x8LX, 9x8RX, 939KS, 939KS/020, 959KS/100, 959KS/200, 959KS/300, 959KS/400, 969KS/100, 969KS/200, 969KS/300, 969KS/400, 969KS/120, 969KS/220, 969KS/320, 969KS/420, 979KS/100, 979KS/200, 979KS/300, 979KS/400, 989/100, 989/200, 989/400, 989/600, 989/150, 989/250, 989/300, 989/350, 989/450, 989/500, 989/550, 990CX 992/100CX, 992/200CX, 992/300CX, 992/400CX, 990DX, 992/100DX, 992/200DX, 992/300DX, 992/400DX, 991CX, 995/100CX, 995/200CX, 995/300CX, 995/400CX, 995/500CX, 995/600CX, 995/700CX, 995/800CX, 991DX, 995/100DX, 995/200DX, 995/300DX, 995/400DX, 995/500DX, 995/600DX, 995/700DX, 995/800DX, 996/80, 996/100, 996/200, 996/300, 996/400, 996/500, 996/600, 996/700, 996/800, 996/900, 996/1000, 996/1200, 997/100, 997/200, 997/300, 997/400, 997/500, 997/600, 997/800	6 months following next Mainline release after 6.5 (about end of year 2000)
* Support life end as of 1/00 ** Support life ends as of 8/00		

Table 8-2. Supported System Release Matrix

Supported Releases	Supported Systems	Support Termination Date
Release 6.5 (65.xx)	<p>9x7, 9x7LX, 9x7RX, 9x7SX, 9x8LX, 9x8RX, 939KS, 939KS/020, 959KS/100, 959KS/200, 959KS/300, 959KS/400, 969KS/100, 969KS/200, 969KS/300, 969KS/400, 969KS/120, 969KS/220, 969KS/320, 969KS/420, 979KS/100, 979KS/200, 979KS/300, 979KS/400, 989/100, 989/200, 989/400, 989/600, 989/150, 989/250, 989/350, 989/300, 989/350, 989/500, 989/550, 989/650</p> <p>990CX 992/100CX, 992/200CX, 992/300CX, 992/400CX, 990DX, 992/100DX, 992/200DX, 992/300DX, 992/400DX, 991CX, 995/100CX, 995/200CX, 995/300CX, 995/400CX, 995/500CX, 995/600CX, 995/700CX, 995/800CX, 991DX, 995/100DX, 995/200DX, 995/300DX, 995/400DX, 995/500DX, 995/600DX, 995/700DX, 995/800DX, 996/80, 996/100, 996/200, 996/300, 996/400, 996/500, 996/600, 996/700, 996/800, 996/900, 996/1000, 996/1200, 997/100, 997/200, 997/300, 997/400, 997/500, 997/600, 997/800, 997/1000, 997/1200</p>	6 months following next Mainline release after 7.0

9 Catalog of User Documentation

This chapter provides listings of customer manuals for the HP 3000 computer system. The listings are divided into two sections:

- “MPE/iX New or Updated Manuals,” which lists all manuals that have been introduced or changed since the MPE/iX 6.0 Release.
- “Manual Collections,” which lists manuals by collections in alphabetical order. For detailed information on a particular manual or manual collection, refer to the *MPE/iX Documentation Guide*.

If your contract includes Material-Based Services, you will receive both software and manual revisions. For additional copies of new or revised manuals, you can order Manual Update Services (MUS).

Many of the learning products listed in this chapter can be individually ordered by calling HP Parts Direct Ordering at 800-227-8164. Specify the “order part number” of the manual you are interested in when ordering.

New to MPE/iX Release 6.5 is the PDF CD-ROM. This CD-ROM contains files for manuals in pdf format. These pdf files are viewable and printable using Adobe Acrobat Reader 3.0 or later. Acrobat Reader also allows you to perform key word searches on the entire CD-ROM contents. If you do not have Acrobat Reader on your system, you can download the 4.0 version from the Adobe website, or install the 3.0 version which is included with the PDF CD-ROM.

You can also view MPE/iX document files on an external web site (<http://www.docs.hp.com>). These files are presented in a html format which is searchable and printable from the web-site.

MPE/iX 6.5 New or Updated Manuals

This section lists customer manuals introduced or updated for MPE/iX 6.5.

Table 9-1. MPE/iX 6.5 New or Updated Manuals

Manual Title	Part No.	Latest Edition
Accessing Files Programmer's Guide	32650-90885	3/00
ALLBASE/SQL Reference Manual	36216-90216	3/00
Configuring and Managing MPE/iX Internet Services	32650-90891	3/00
Communicator 3000 MPE/iX Release 6.5	30216-90291	3/00
HP 3000 MPE/iX System Software Maintenance Manual	30216-90298	3/00
HP ARPA File Transfer Protocol User's Guide	36957-90159	3/00
HP Data Entry and Forms Management System (VPLUS)	32209-90024	3/00
Legato NetWorker ClientPak and Storage Node for MPE/iX Installation Guide	B6266-90006	8/99
Legato NetWorker ClientPak and Storage Node for MPE/iX Supplement	B6266-90004	12/98
Legato NetWorker ClientPak and Storage Node for MPE/iX User's Guide	B6266-90005	8/99
MPE/iX Commands Reference Manual	32650-90877	3/00
MPE/iX Documentation Guide	32650-90890	3/00
MPE/iX Error Messages Volume 1	32650-90878	3/00
MPE/iX Error Messages Volume 2	32650-90879	3/00
MPE/iX Error Messages Volume 3	32650-90880	3/00
MPE/iX Glossary of Terms and Acronyms	32650-90887	3/00
MPE/iX Intrinsic Reference Manual	32650-90875	3/00
MPE/iX Quick Reference Guide	32650-90881	3/00
MPE/iX System Utilities Reference Manual	32650-90882	3/00
ODBCLink/SE Reference Manual	36217-90409	8/99
Query Reference Manual for MPE/iX	32650-90889	3/00
Sort/Merge/XL General Programmer's Guide	32650-90884	3/00
Sort/Merge/XL General User's Guide	32650-90883	3/00
STORE and TurboSTORE/iX Products Manual	B5151-90003	3/00

Table 9-1. MPE/iX 6.5 New or Updated Manuals

Manual Title	Part No.	Latest Edition
System Debug Manual	32650-90888	3/00
TurboIMAGE/XL Database Management System Reference Manual	30391-90011	3/00
Using KSAM/XL	32650-90886	3/00

Manual Collections

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
Communicators and System Software Maintenance Manuals					
Communicator 3000 (6.5)	30216-90291	30216-90291	Mar-00	B	y
Communicator 3000 (PowerPatch 1 based on 6.0)	30216-90286	30216-90286	Aug-99	B	y
Communicator 3000 (6.0)	30216-90269	30216-90269	Oct-98	B	y
Communicator 3000 (PowerPatch 5 based on 5.5)	30216-90257	30216090257	Aug-98	P	y
Communicator 3000 (PowerPatch 7 based on 5.5)	30216-90282	30216-90282	Mar-99	P	y
HP 3000 MPE/iX System Software Maintenance Manual (6.5)	30126-90298	30216-90298	Mar-00	B	y
HP 3000 MPE/iX System Software Maintenance Manual (6.0)	30126-90272	30216-90272	Oct-98	B	y
HP 3000 MPE/iX System Software Maintenance Manual (5.5)	30216-90223	30216-90223	Jul-96	P	n
MPE/iX Operating System					
CI Programming Quick Reference Pocket Card	32650-90316	32650-90269	Mar-00	HC	n
FCOPY Reference Manual	32212-90003	32212-90008	Jun-92	B	y
HP 3000 Series 9X8LX Computer Systems: Commands Reference	B3813-90011	B3813-90012	Apr-94	B	y
HP 3000 Series 9X8LX Computer Systems: Getting Started	B3813-90003	B3813-90014	Apr-94	B	y
HP 3000 Series 9X8LX Computer Systems: Task Reference Manual	B3813-90009	B3813-90010	Apr-94	B	y
HP 3000 Series 9X8LX Computer Systems: Understanding Your System	B3813-90001	B3813-90013	Apr-94	B	y
HP Easytime/XL Quick Reference Card	B1940-90001	B1940-90611	Jun-92	HC	n
* The documentation website address is: http://www.docs.hp.com .					

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
HP Easytime/XL User's Guide	B1940-90002	B1940-90602	Sep-91	HC	n
MPE/iX Commands Reference Manual	32650-90877	32650-90877	Mar-00	B	y
MPE/iX Day to Day Tasks Reference Card	A1707-90004	A1707-96034	Jun-92	HC	n
MPE/iX Documentation Guide	32650-90890	32650-90890	Mar-00	B	y
MPE/iX Error Messages Manual Vol 1	32650-90878	32650-90878	Mar-00	B	y
MPE/iX Error Messages Manual Vol 2	32650-90879	32650-90879	Mar-00	B	y
MPE/iX Error Messages Manual Vol 3	32650-90880	32650-90880	Mar-00	B	y
MPE/iX Glossary of Terms and Acronyms	32650-90887	32650-90887	Mar-00	B	y
MPE/iX Quick Reference Guide	32650-90881	32650-90881	Mar-00	B	y
MPE/iX System Utilities Reference Manual	32650-90882	32650-90882	Mar-00	B	y
SORT-MERGE/XL General User's Guide	32650-90883	32650-90883	Mar-00	B	y
SORT-MERGE/XL Programmer's Guide	32650-90884	32650-90884	Mar-00	B	y
Using HP 3000: Advanced Skills Tutorial	32650-90872	32650-90872	Oct-98	B	y
Using HP 3000: Fundamental Skills Tutorial	32650-90871	32650-90871	Oct-98	B	y
Development Tools and Distributed Computing					
Accessing Files Programmer's Guide	32650-90885	32650-90885	Mar-00	B	y
ALLBASE NET User's Guide	36216-90031	36216-90101	Apr-94	B	y
ALLBASE/SQL Advanced Application Programming Guide	36216-90100	36216-90099	Apr-94	B	y
ALLBASE/SQL C Application Programming Guide	36216-90023	36216-90080	Jun-92	B	y
ALLBASE/SQL COBOL Application Programming Guide	36216-90006	36216-90081	Jun-92	B	y

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
ALLBASE/SQL Database Administration Guide	36216-90005	36216-90214	Aug-97	B	y
ALLBASE/SQL FORTRAN Application Programming Guide	36216-90030	36216-90079	Jun-92	B	y
ALLBASE/SQL Message Manual	36216090213	36216-90009	Aug-97	B	y
ALLBASE/SQL Pascal Application Programming Guide	36216-90007	36216-90082	Oct-92	B	y
ALLBASE/SQL Performance and Monitoring Guidelines	36216-90102	36216-90103	Apr-94	B	y
ALLBASE/SQL Reference Manual	36126-90216	36216-90216	Mar-00	B	y
Asynchronous Serial Communications Programmer's Reference Manual	32022-61001	32022-90052	Oct-98	B	y
BASIC Interpreter Reference Manual	30000-90026	30000-90026	Nov-78	HC	n
BASIC/3000 Compiler Reference Manual	32103-90001	32103-90001	Sep-77	HC	n
Berkeley Sockets/iX Reference Guide	32650-90363	32650-90372	Apr-94	B	y
Command Interpreter Access and Variables Programmer's Guide	32650-90011	32650-90493	Apr-94	B	y
Compiler Library/XL Reference Manual	32650-60014	32650-90029	Oct-88	HC	y
Data Dictionary Managing Information Network Primer	5958-8527	5958-8527	Nov-86	HC	n
Data Types Conversion Programmer's Guide	32650-60010	32650-90015	Oct-89	B	y
DBChange Plus Technical Addendum for MPE/iX Release 4.0	36386-90005	36386-90005	Jun-92	P	y
DBChange Plus User's Guide	36386-90001	36386-90003	Feb-91	P	y
DCE for the HP 3000	B3821-90001	B3821-90002	Oct-95	HC	y
Dictionary/3000 Reference Manual	32244-90001	32244-90001	Dec-87	HC	y
Dictionary/3000 Documentation Update Notice	32244-90013	32244-90013	Oct-96	HC	n

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
EDIT/3000 Reference Manual	03000-90012	32650-90385	Aug-80	B	y
Getting Started as an MPE/iX Programmer	32650-90008	32650-90421	Jun-92	B	y
Getting Started with HP IMAGE/SQL	36385-90008	36385-90011	Dec-94	B	y
Getting Started with HP Software Revision Controller (SRC)	30234-60002	30234-90002	Nov-88	HC	n
Getting Started with TRANSACT	32247-60002	32247-90007	Jul-88	HC	y
High Level Screen Management Intrinsic Library Reference Manual	32424-60001	32424-90002	Nov-87	HC	y
HP 3000 Basic for Beginners	03000-90025	03000-90025	Nov-72	HC	n
HP ALLBASE Query User Guide Kit	32426-64001	32426-64001	Oct-89	HC	n
HP ALLBASE/4GL Developer Administration Manual	30601-90201	30601-90205	May-92	P	y
HP ALLBASE/4GL Developer Quick Reference Guide	30601-90210	30601-90211	May-92	P	y
HP ALLBASE/4GL Developer Reference Manual Vol 1	30601-90202	30601-90206	May-92	P	y
HP ALLBASE/4GL Developer Reference Manual Vol 2	30601-90204	30601-90208	May-92	P	y
HP ALLBASE/4GL Installation Manual	30601-64212	30601-64213	May-92	HC	n
HP ALLBASE/4GL Developer Self-Paced Training Guide	30601-90203	30601-90207	May-92	P	y
HP ALLBASE/4GL Software Update Notice B.06 Release	5961-7797	5063-3725	Feb-93	HC	y
HP ALLBASE/BRW Reference Manual	35360-90051	35360-90052	Jan-92	P	y
HP ALLBASE/BRW SW Update Notice for MPE/iX (BRW A.01.50)	35360-90204	35360-90203	Apr-94	P	y
HP ALLBASE/BRW Tutorial	35360-90201	35360-90202	May-92	P	y
HP Browse/XL User's Guide	36384-90001	36384-90001	Oct-90	HC	y
HP Business BASIC Programmer's Guide	32115-90003	32115-90003	Jul-87	HC	n

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
HP Business BASIC Quick Reference Guide	32115-90002	32115-90002	Jul-85	HC	n
HP Business BASIC Reference Manual	32115-90001	32115-90001	Jul-87	HC	n
HP Business BASIC/XL Migration Guide	32715-60002	32715-90003	Oct-89	HC	y
HP Business BASIC/XL Reference Manual	32715-60001	32715-90001	Oct-89	HC	y
HP C Programmer's Guide	92434-90002	92434-90009	Jul-96	P	y
HP C/iX Library Reference Manual	30026-90001	30026-90004	Oct-92	P	y
HP C/iX Reference Manual	31506-90005	31506-90011	Jun-92	P	y
HP COBOL II/V Reference Manual	32233-90001	32233-90001	May-89	HC	n
HP COBOL II/XL Migration Guide	31502-60011	31500-90004	Oct-88	HC	y
HP COBOL II/XL Programmer's Guide	31500-90002	31500-90014	Jul-91	P	y
HP COBOL II/XL Quick Reference Guide	31500-90003	31500-90015	Jul-91	HC	y
HP COBOL II/XL Reference Manual	31500-90001	31500-90013	Jul-91	P	y
HP Data Entry and Forms Mgmt System (VPLUS) Reference Manual	32209-90024	32209-90024	Mar-00	B	y
HP EDIT Quick Reference Guide	30316-90005	30316-90017	Dec-90	P	n
HP EDIT Reference Manual	30316-90001	30316-90016	Dec-90	P	y
HP FORTRAN 77/iX Migration Guide	31501-90004	31501-90023	Jun-92	P	y
HP FORTRAN 77/iX Programmer's Guide	31501-90011	31501-90022	Jun-92	P	y
HP FORTRAN 77/iX Reference Manual	31501-90010	31501-90021	Jun-92	P	y
HP GlancePlus User's Manual (for MPE/iX Systems)	B1787-90001	B1787-90008	Apr-92	HC	y
HP IMAGE/SQL Administration Guide	36385-90001	36385-90012	Aug-97	B	y

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
HP Link Editor/iX Reference Manual	32650-90030	32650-90309	Dec-90	P	y
HP Link Editor/iX Technical Addendum	32650-09476	32650-90845	Oct-95	P	y
HP Pascal/iX Migration Guide	31502-60011	31502-90004	Nov-87	HC	y
HP Pascal/iX Programmer's Guide	31502-90002	31502-90023	Jun-92	HC	y
HP Pascal/iX Reference Manual	31502-90001	31502-90022	Jun-92	HC	y
HP Performance Collection Software User's Manual	50700-90022	50700-90038	Apr-92	P	n
HP RPG/iX Pocket Guide	30318-90002	30318-90002	Oct-89	HC	n
HP RPG/iX Programmer's Guide	30318-60001	30318-90001	Jul-89	HC	y
HP RPG/iX Reference Manual	30318-60002	30318-90011	Dec-93	P	y
HP RPG/iX Reference Manual Software Update Notice	30318-90016	30318-90017	Mar-95	P	y
HP RPG/iX Utilities Reference Manual	30318-60003	30318-90006	Oct-89	HC	y
HP Search/XL User's Guide	36383-90001	36383-90001	Oct-90	HC	y
HP Software Revision Controller (SRC) Implementation Guide	30234-60002	30234-90003	Nov-88	HC	n
HP Software Revision Controller (SRC) Quick Reference Card	30234-60002	30234-90005	Nov-88	HC	n
HP Software Revision Controller (SRC) User's Guide	30234-60002	30234-90001	Nov-88	HC	n
HP Software Revision Controller/XL Product Information Update	30234-60002	30234-90006	Nov-88	HC	n
HP Symbolic Debugger/iX User's Guide	31508-90003	31508-90014	Jun-92	P	y
HP System Dictionary/XL COBOL Definition Extractor	32257-90001	32257-90001	Dec-87	HC	y
HP System Dictionary/XL General Reference Manual Vol 1	32256-90004	32256-90004	May-88	HC	y
HP System Dictionary/XL General Reference Manual Vol 2	32256-90005	32256-90005	May-88	HC	y

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
HP System Dictionary/XL Intrinsic Reference Manual	32256-90002	32256-90002	May-88	HC	y
HP System Dictionary/XL SDMAIN Reference Manual	32256-90001	32256-90001	May-88	HC	y
HP System Dictionary/XL Self-Paced Customer Training	32254-91001	32254-91001	Aug-87	HC	n
HP System Dictionary/XL Utilities Reference Manual	32256-90003	32256-90003	May-88	HC	y
HP Toolset/XL Reference Manual	36044-60001	36044-90001	Jan-84	HC	y
HP TRANSACT Documentation Update Notice	32247-90028	32247-90028	Oct-96	P	y
HP TRANSACT Quick Reference Guide	32247-90020	32247-90027	Oct-96	P	y
HP TRANSACT Reference Manual	32247-60003	32247-90026	Apr-94	P	y
INFORM/V User's Guide	32246-60002	32246-60002	Mar-88	HC	y
Introduction to MPE/XL for MPE V Programmers	30367-90005	30367-90005	Oct-89	P	y
Introduction to MPE/XL for MPE V System Administrators	30367-90003	30367-90017	Dec-90	P	y
IPC Communications Programmer's Guide	32650-90019	32650-90019	Nov-87	B	y
ISQL Reference Manual for ALLBASE/SQL & IMAGE/SQL	36216-90096	36216-90095	Apr-94	B	y
KSAM/3000 Reference Manual	30000-90079	32650-90386	Jun-92	B	y
Learning HP EDIT	30316-90002	30316-90015	Dec-90	P	y
Message Catalogs Programmer's Guide	32650-90021	32650-90021	Mar-90	B	y
Migration Process Guide	30367-90007	30367-90019	Jun-92	P	y
MPE Segmenter Reference Manual	30000-90011	30000-90011	Aug-86	P	y
MPE V to MPE XL: Getting Started Mentor's Guide	30367-90004	30367-90004	Oct-89	P	y
MPE V to MPE XL: Getting Started Self-Paced Training	30367-90002	30367-90002	Oct-89	P	y

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
MPE/iX AIF: OS Reference Manual	36374-90001	36374-90013	Dec-94	B	y
MPE/iX Developer's Kit Reference Manual Vol 1	36430-90001	36430-90007	Apr-94	B	y
MPE/iX Developer's Kit Reference Manual Vol 2	36430-90002	36430-90008	Apr-94	B	y
MPE/iX Intrinsic Reference Manual	32650-90875	32650-90875	Mar-00	B	y
MPE/iX Shell and Utilities Reference Manual Vol 1	36431-90001	36431-90007	Apr-94	P	.pdf
MPE/iX Shell and Utilities Reference Manual Vol 2	36431-90003	36431-90008	Apr-94	P	.pdf
MPE/iX Shell and Utilities User's Guide	36431-90002	36431-90006	Apr-94	HC	n
Native Language Programmer's Guide	32650-90022	32650-90207	Sep-91	B	y
ODBCLINK/SE Reference Manual	36217-90406	36217-90409	Aug-99	P	y
Process Management Programmer's Guide	32650-90023	32650-90023	Nov-87	B	y
QUERY Reference Manual for MPE/iX	30000-90889	32650-90889	Mar-00	B	y
Resource Management Programmer's Guide	32650-90024	32650-90024	Nov-87	B	y
RPG/V Reference Manual	32104-90001	32104-90001	Sep-89	HC	n
SPL to HP C/XL Migration Guide	30231-60001	30231-90001	Oct-89	HC	y
Switch Programming Guide	32650-60030	32650-90014	Nov-87	P	y
System Debug Reference Manual	32650-90888	32650-90888	Mar-00	B	y
The POSIX.1 Standard: A Programmer's Guide ISBN-0-8053-9605-5	-	36430-90006	Apr-94	HC	n
Trap Handling Programmer's Guide	32650-90026	32650-90026	Mar-90	P	y
TurboIMAGE/XL Database Management System Reference Manual	30391-90011	30391-90011	Mar-00	B	y

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
Up and Running with ALLBASE/SQL	36389-90011	36389-90016	Jun-92	B	y
User Logging Programmer's Guide	32650-90027	32650-90027	Jul-88	B	y
Using KSAM/XL	32650-90886	32650-90886	Mar-00	B	y
Using VPLUS/V: An Introduction to Forms Design	32209-90004	32209-90004	Aug-86	HC	y
High Availability					
Auto/Restart/XL User's Guide	36375-90001	36375-90004	Oct-92	B	y
Installation/Upgrade Procedures for SHAREPLEX	B3933-90005	B3933-90005	Jun-97	P	.pdf
Legato NetWorker ClientPak for MPE/iX Installation Guide	B5475-90001	B5475-90001	Dec-97	P	.pdf
Legato NetWorker ClientPak for MPE/iX Release 5.0 Supplement	B5475-90005	B5475-90005	Feb-98	P	.pdf
Legato NetWorker ClientPak for MPE/iX Users Guide	B5475-90002	B5475-90002	Dec-97	P	.pdf
Legato NetWorker ClientPak and Storage Node for MPE/iX Installation Guide	B6266-90006	B6266-90006	Aug-99	P	.pdf
Legato NetWorker ClientPak and Storage Node for MPE/iX Users Guide	B6266-90005	B6266-90005	Aug-99	P	.pdf
Legato NetWorker ClientPak and Storage Node for MPE/iX Release 5.5 Supplement	B6266-90004	B6266-90004	Dec-98	P	.pdf
Mirrored Disk/iX User's Guide	30349-90003	30349-90005	Jun-92	B	y
NBSpool for MPE Reference Guide	B3933-90002	B3933-90002	Jun-97	P	.pdf
NBSpool Release Notes Version 9.7	B3933-90008	B3933-90008	Jun-97	P	.pdf
NetBase for MPE Reference Guide	B3933-90001	B3933-90001	Jun-97	P	.pdf
NetBase Release Notes Version 9.7	B3933-90007	B3933-90007	Jun-97	P	.pdf
NetBase SQL Shadowing Handbook	B3482-90006	B3482-90006	Jun-97	P	.pdf

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
VISTA Plus for MPE Administrators Guide	B3933-90004	B3933-90004	Jun-97	P	.pdf
VISTA Plus for MPE User's Guide	B3933-90003	B3933-90003	Jun-97	P	.pdf
VISTA Release Notes Version A.07	B3933-90006	B3933-90006	Jun-97	P	.pdf
System Hardware and Diagnostics					
CPU Upgrade Manual (9x9KS)	-	A2375-90012	Jan-98	P	.pdf
Expansion Cabinet Installation Guide (99X/Family)	-	A1809-90006	May-97	P	.pdf
Firmware Update Guide (99X/Family)	-	A1820-90002	May-97	P	.pdf
Firmware Update Quick Reference (99X/T-Class)	-	A1820-90003	May-97	P	.pdf
HP 3000 9X9KS Site Preparation and Requirements Guide	-	A2375-90073	Nov-96	P	.pdf
HP 3000 and HP 9000 PA-RISC Computer System Support Log	09740-90013	09740-96033	Feb-96	P	.pdf
HP 3000 CS 99x/890/T500 Families Operator's Guide	A1809-90009	A1809-96019	Apr-97	P	n
HP 3000 Series 9X8LX/RX Computer Systems Inst. and Configuration Guide	A2051-90006	A2051-96006	Oct-93	P	n
HP 3000/9x9KS Installation Guide	-	-	-	P	.pdf
HP PA-RISC Computer Systems Integrated Cabinet Installation Guide (9X9KS)	-	A2375-90007	Jan-98	P	.pdf
HP PA-RISC Computer Systems System Upgrade Guide (9X9KS)	-	A2375-90010	Feb-97	P	.pdf
I/O Upgrade Manual (9X9KS)	-	A2375-90027	Jan-98	P	.pdf
Installation Guide (99X Family)	-	A1809-90001	Oct-97	P	.pdf
Internal Peripherals Update Guide (9X9KS)	-	A2375-90008	Jan-98	P	.pdf
Operator's Guide (99X Family)	-	A1809-90009	Jun-97	P	.pdf

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
Site Preparation and Requirements Guide (99X Family)	-	A1809-90002	May-97	P	.pdf
System Memory Upgrade Guide (9X9KS)	-	A2375-90009	Jan-98	P	.pdf
System Upgrade Installation Guide (CS 99X Family) (990/992 to 996)	-	A3310-90002	May-97	P	.pdf
System Upgrade Installation Guide (CS 99X Family) (991/995 to 996)	-	A3310-90001	May-97	P	.pdf
Networking and Communications					
APPC Subsystem on MPE/XL Node Manager's Guide	30294-61002	30294-90007	Jun-92	B	y
Configuring and Managing Host-Based X.25 Links	36939-61004	36939-90054	Oct-98	B	y
HP ARPA File Transfer Protocol User's Guide	36957-90159	36957-90159	Mar-00	B	y
HP SNA Products Remote System Configuration Guide	J2220-61025	30292-90008	Mar-95	B	y
HP SNMP/XL User's Guide	36922-61029	36922-90036	Mar-94	B	y
HP Telnet/iX User's Guide	36957-90154	36957-90156	Oct-98	B	y
HP-PB 100Base-T Network Adapter Installation and Service Guide	B5427-90001	B5427-90001	Aug-97	B	y
HP-PB 100VG-AnyLAN Network Adapter Installation and Service Guide	B5426-90001	B5426-90001	Aug-97	B	y
LU 6.2 API Application Programmer's Reference Guide	30294-61000	30294-90008	Jun-92	B	y
Managing Host-Based X.25 Links Quick Reference Card	36939-61003	36939-90051	Dec-94	B	y
NetIPC 3000/XL Programmer's Reference Manual	36920-61005	5958-8600	Oct-89	B	y
NS Cross-System NFT Reference Manual	36920-61003	5960-1634	Jan-89	HC	n
NS3000/iX Error Messages Reference Manual	36923-61000	36923-90041	Oct-98	B	y

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
NS3000/iX NMMGR Screens Reference Manual	36922-61003	36922-90038	Oct-98	B	y
SNA IMF Programmer's Reference Manual	30293-61005	30293-90009	Jun-92	B	y
SNA IMF/XL Node Manager's Guide	30293-61000	30293-90010	Jun-92	B	y
SNA Link/iX Node Manager's Guide	30291-61000	30291-90009	Mar-94	B	y
SNA NRJE Node Manager's Guide	30292-61000	30292-90007	Oct-92	B	y
SNA NRJE User/Programmer Reference Manual	30292-61001	30292-90006	Oct-92	B	y
Using NS3000/iX Network Services	36920-61000	36920-90008	May-94	B	y
Using SNA IMF Pass Thru	30293-61008	30293-90006	Dec-90	B	y
Network and System Management					
Configuring and Managing MPE/iX Internet Services	32650-90891	32650-90891	Mar-00	B	y
Configuring Systems for Terminals, Printers, and Other Serial Devices	32022-61000	32022-90051	Oct-98	B	y
Customizing Terminal and Printer Type Files with the Workstation Configurator	5959-2870	32022-90031	Feb-94	B	y
HP 3000/iX Network Planning and Configuration Guide	36922-61023	36922-90037	Oct-98	B	y
HP OpenView Console Manager's Guide	B3118-90002	B3118-90012	Aug-92	B	y
HP OpenView Console User's Guide	B3118-90001	B3118-90011	Aug-92	B	y
HP Security Monitor/iX Manager's Guide	32650-90455	32650-90498	Apr-94	B	y
HP Security Monitor/iX User's Guide	32650-90454	32650-90497	Apr-94	B	y
Installing and Managing HP ARPA File Transfer Protocol Network Manager's Guide	36957-90159	36957-90159	Oct-98	B	y

Table 9-2. Manual Collections

Manual Title	Order Part Number	Part Number	Latest Edition	How Available? HC = Hardcopy only P = PDF CD B = PDF & Inst. Info	View on the WEB*
Manager's Guide to MPE/iX Security	32650-90474	32650-90473	Apr-94	B	y
Managing Spooler Operations Quick Reference Pocket Card	32650-90268	32650-90488	Apr-94	HC	n
MPE/iX HP 3000 Series 99X Software Startup Manual	36123-90046	36123-90046	Apr-94	P	n
Native Mode Spooler Reference Manual	32650-90166	32650-90867	Oct-98	B	y
New Features of MPE/iX: Using the Hierarchical File System	32650-90351	32650-90492	Apr-94	B	y
NS3000/iX Operations and Maintenance Reference Manual	36922-61005	36922-90039	Oct-98	B	y
Openview DTC Technical Reference Manual	5961-9820	5961-9882	Oct-94	HC	n
Performing System Management Tasks	32650-90004	32650-90854	Jul-96	B	y
Performing System Operation Tasks	32650-90137	32650-90484	Apr-94	B	y
STORE and TurboSTORE/iX Products Manual	B5151-90003	B5151-90003	Mar-00	B	y
System Startup, Configuration, and Shutdown Reference Manual	32650-90042	32650-90855	Jul-96	B	y
Troubleshooting Terminal, Printer and other Serial Devices	32022-61002	32022-90030	Oct-93	B	y
User's Guide to MPE/iX Security	32650-90472	32650-90471	Apr-94	B	y
Using HP OpenView DTC Manager	D2355-90001	D2355-95018	Jan-93	HC	n
Using the HP 3000 Workload Manager	B3879-90001	B3879-90002	Dec-94	B	y
Using the Node Management Services (NMS) Utilities	32022-61005	32022-90053	Oct-98	B	y
Volume Management Reference Manual	32650-90045	32650-90491	Apr-94	B	y
* The documentation website address is: http://www.docs.hp.com .					

Numerics

12H

configuration, 115

A

ABORTJOB command, 153

acceptable SQL, 158

AIF:OS

error messages, 56

AIFFILEGGET, 39, 40, 55

AIFFILEGPUT, 40, 55

AIFFILELGET, 39, 55

AIFFILELPUT, 39, 55

AIFJSGET, 57

AIFSYSWIDEGET, 57

ALLBASE/SQL database, 148

ALTJOB, 66, 71

Apache

components, 139

configuring, 140

documentation, 144

error logging, 143

features, 138

starting, 139

stopping, 143

verify activity, 142

Apache for MPE/iX, 138

applet, 149

Architected interfaces, 39

archive file, 150

B

Backdating, 53

B-TREE, 132

BUILD command, 41

C

Checkpoint, 131

changing, 131

controlling, 131

child processes, 152

class files, 153

CLASSPATH, 153

Client Components

Java, 150

client components

Java class files, 150

sample source file, 150

client files, 150

client tracing, 159

client-server application, 148

Command Interpreter, 74

Command interpreter, 41

Commands

BUILD, 41

NEWCI, 74

RUN, 74

commands

FILE, 44

LISTFILE, 45

Commercial Systems MPE/iX releases and products, 173

Compatibiity mode, 64

components, 149

Concurrent processes, 74

Configuring

disks, 63

configuration option

LOGFILE, 151

TIMEOUT, 151

connection parameters, 154

Console messages, 96

Console Messages template, 93

conversion grid, 157

D

Data Center Management, 88

Data segments, 64

data type conversions, 156

data type mapping, 156

Database monitoring, 96

Database Monitoring template, 93

databases, 149

DBMS-independent protocol, 148

DCM, 88

decimal value, 157

default port number 31700, 152

default server configuration file, 151

discfree, 63

Disgnostic platform, 91

Disks, 63

Dispatcher

Multiple runqueues, 60

DLT drives, 101

configuration, 106

HW requirements, 103

performance, 104

SW requirements, 103

DLT4000 differential drives, 103

DLT7000 differential drives, 103

DLT8000 differntial drive, 107

documentation list, 181

driver components

Java archive (JAR), 150
UNIX tape archive (TAR), 150
Windows ZIP (ZIP), 150
Driver Manager, 154
dropped connections, 153
dstat, 63
DSTs, 64

E

Electronic Support Center, 23
Enterprise Management Solution, 88
ENVIRONMENT VARIABLE, 153
expanded B-tree, 132
extracting driver class files, 151
 Win32 Platform, 151
extracting file, 150

F

fast search, 101
Fast search on tapes, 101
FFILEINFO, 38
FGETINFO, 38
File Open, 135
file open, 132
FLABELINFO, 37, 38, 47
floating-point value, 157

H

HiPRI, 66
HP Driver for JDBC, 148, 149
HP Driver for JDBC components, 150
HP Electronic Support Center, 29
HP JDBC Client components, 149
HP JDBC Driver, 153
HP JDBC Monitor, 151
HP JDBC Server components, 151
HP Secure Web Console, 89
 installation, 98
 product features, 99
HP SupportLine
 MPE/iX patches available, 29
HPFADDTOPINTER, 48
HPFFILLDATA, 52
HPFMOVEDATA, 49
HPFMOVEDATALTOR, 50
HPFMOVEDATARTOL, 51
HPFOPEN, 47
HP-IB device adapter, 12
HP-UX server components, 150

I

IMAGE database, 53
IMAGE/SQL database, 148
INSERT statement, 158
installation script, 150
integer value, 157
Internet
 HP SupportLine, 29
 patches available, 29

J

JAR, 150
Java 1.2, 146
Java API, 148
Java application, 148
Java class path, 153
Java Client Components, 150
Java compiler, 153
Java components, 149
Java Requirements, 149
Java Run Time Environment, 153
Java SDK, 148, 154
Java Virtual Machine, 154
Java Virtual Machines, 148
Java VM rules, 157
JAVA/iX, 151
JDBC
 Driver Manager, 148
 Java application, 148
JDBC API, 148, 149
JDBC client-server application, 148
JDBC Driver, 149
JDBC Driver Manager, 148, 154
JDBC Monitor, 149, 152
 shutdown, 152
 startup, 152
JDBC Server, 149
JDBC server components, 150
JDBC SQL, 149
JDK, 149, 154
JDK components, 149
JOB, 70
Job queues
 user-defined, 66
JSTOPMON stream job, 153
JSTRTMON stream job, 153

K

kill command, 152
killing the monitor, 152

L

Large file

- 32-bit interger overflow, 37
- AIF:OS impact, 39
- backdating, 53
- CI changes, 41
- compatibility, 35
- data access, 36
- FILE command, 44
- functionality, 34
- IMAGE database, 53
- intrinsic changes, 46, 47
- LISTFILE command, 45
- space ID boundary, 36
- STORE/RESTORE impact, 52
- support, 34
- user-mapped access, 35
- utility and subsys support, 53

Large files, 34

Large transactions

- limitations, 130
- with IMAGE, 126
- with transaction manager, 127
- with TurboIMAGE, 128

LDAP C-SDK/iX, 145

learning products

- ordering, 171

Legato backup solution, 88

LIMIT, 69

LISTJOBQ, 68

listjobq, 66

loading HP JDBC Driver, 154

log file, 151

LOGFILE, 151

logging levels, 152

M

mainline release, 176

manual list, 181

marker character, 158

MESA, 110

method arguments, 155

mirvutil, 63

monctrl command, 152

monctrl tool, 152

MPE/iX, 151

- releases and products, 173

MPE/iX HP JDBC Monitor, 152

MPE/iX Monitor shutdown, 152

MPE/iX Monitor startup, 152

MPE/iX Server Components, 150

Multiple job queues, 57

Multiple runqueues, 60

N

Network monitoring, 95

NEWCI command, 74

CI.PUB.SYS, 77

logon UDC, 77

parameters, 76

progfile, 75

session resources, 76

syntax, 76

NEWJOBQ, 67

newjobq, 66

NNMGR, 58

non-platform release, 176

numeric casting, 158

O

Obsolete products, 14

Omniback agent, 89

Open View system manager, 97

ordering

PowerPatch 5, 171

products, 171

OS Monitoring template, 93

P

parameterized SQL statments, 158

patches

on Internet, 29

platform release, 176

platforms

HP-UX, 151

MPE/iX, 151

portnumber parameter, 152

PowerPatch 5

ordering, 171

Priority control, 131

Procedures, 159

product

changes by releases, 173

product list

how to order, 171

product release history, 173

PURGEJOBQ, 68

purgejobq, 66

R

releases

mainline, 176

- non-platform, 176
- platform, 176
- product history, 173
- supported, 176
- requirements, 149
- RUN command, 74

S

- sample client, 153, 154
- Security messages, 96
- SELECT statement, 158
- servcfg, 151
- Server Components
 - MPE/iX, 150
- server components, 150
- server configuration file, 160
- server log file, 152
- Server Logging, 152
- server logging, 160
- server logging levels, 160
- Server Requirements
 - ALLBASE/SQL, 150
 - HP-UX, 149
 - IMAGE/SQL, 150
 - MPE/iX, 150
- SHOWJOB, 73
- shutdown, 152
- simple client, 153
- SimpleClient, 153
- space characters, 160
- Spindles, 63
- Starting HP JDBC Monitor, 152
- startup, 152
- startup script, 150
- STM, 91
- Stopping HP JDBC Monitor, 152
- stored procedures, 159
- STREAM, 72
- stream jobs, 152
 - JSTOPMON, 152
 - JSTRTMON, 152
- stream script, 150
- suggested conversions, 156
- Sun-compliant JDK version 1.1, 149
- Support tools manager, 91
- supported releases, 176
- SupportLine
 - MPE/iX patches available, 29
 - via Internet, 29
- System limits, 62

T

- TAR, 150
- tar file, 150
- TCP connections
 - configure through NMMGR, 58
- Template
 - console messages, 93
 - database monitoring, 93
 - OS monitoring, 93
- Templates
 - OS monitoring, 93
- temporary directory, 150
- TERMDSM, 110
- TIMEOUT, 151
- tracing, 159
- tracing URL syntax, 159
- Transact/iX, 132
- TurboSTORE, 101

U

- UNIX kill command, 152
- UNIX tape archive
 - TAR, 150
- Unsupported ALLBASE/SQL Data Types
 - BINARY, 158
 - LONG BINARY, 158
 - LONG VARBINARY, 158
- unsupported SQL statements, 158
- URL Syntax, 154
- user
 - documentation list, 181
- User-defined job queues, 66
- Utility and subsys support, 53

V

- value of 0, 157
- value of 1, 157
- Volume management, 96
- volutil, 63

Z

- ZIP, 150