```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

Copyright (c) 2010 CA. All rights reserved.

This documentation (the "Documentation") and related computer software program (the "Software") (hereinafter collectively referred to as the "Product") is for the end user's informational purposes only and is subject to change or withdrawal by CA at any time.

This Product may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Product is proprietary information of CA and protected by the copyright laws of the United States and international treaties.

Notwithstanding the foregoing, licensed users may print a reasonable number of copies of the Documentation for their own internal use, and may make one copy of the Software as reasonably required for back-up and disaster recovery purposes, provided that all CA copyright notices and legends are affixed to each reproduced copy. Only authorized employees, consultants, or agents of the user who are bound by the provisions of the license for the Software are permitted to have access to such copies.

This right to print copies of the Documentation and to make a copy of the Software is limited to the period during which the license for the Product remains in full force and effect. Should the license terminate for any reason, it shall be the user's responsibility to certify in writing to CA that all copies and partial copies of the Product have been returned to CA or destroyed.

EXCEPT AS OTHERWISE STATED IN THE APPLICABLE LICENSE AGREEMENT, TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS PRODUCT "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO THE END USER OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS PRODUCT, INCLUDING WITHOUT LIMITATION, LOST PROFITS, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED OF SUCH LOSS OR DAMAGE.

The use of any product referenced in the Documentation is governed by the end user's applicable license agreement.

The manufacturer of this Product is CA.

This Product is provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Section 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7013(c)(1)(ii), as applicable, or their successors.

All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

CA MICS RESOURCE MANAGEMENT PRODUCT NAMES

Over the years our products' names have changed. The following table shows the most often used names.

PRODUCT'S OLD NAME	PRODUCT'S NEW FULL NAME	PRODUCT'S NEW SHORT NAME
NeuMICS Accounting and Chargeback	CA MICS Resource Management Accounting and Chargeback Option	CA MICS Accounting and Chargeback Option
NeuMICS IDMS Analyzer	CA MICS Resource Management Analyzer Option for CA IDMS	CA MICS Analyzer Option for CA IDMS
NeuMICS TSO Analyzer	CA MICS Resource Management Analyzer for TSO	CA MICS Analyzer for TSO
NeuMICS CICS Analyzer	CA MICS Resource Management Analyzer Option for CICS	CA MICS Analyzer Option for CICS
NeuMICS DB2 Analyzer	CA MICS Resource Management Analyzer Option for DB2	CA MICS Analyzer Option for DB2
NeuMICS IMS Analyzer	CA MICS Resource Management Analyzer Option for IMS	CA MICS Analyzer Option for IMS
NeuMICS Analyzer for MeasureWare OS Agents	CA MICS Resource Management Analyzer Option for Measureware	CA MICS Analyzer Option for MeasureWare
NeuMICS MQSeries Analyzer	CA MICS Resource Management Analyzer Option for MQSeries	CA MICS Analyzer Option for MQSeries
NeuMICS VAX/VMS Analyzer	CA MICS Resource Management Analyzer Option for VAX/VMS	CA MICS Analyzer Option for VAX/VMS

PRODUCT'S OLD NAME	PRODUCT'S NEW FULL NAME	PRODUCT'S NEW SHORT NAME
NeuMICS VM and CMS Analyzer	CA MICS Resource Management Analyzer Option for VM/CMS	CA MICS Analyzer Option for VM/CMS
NeuMICS VSE/POWER Analyzer	CA MICS Resource Management Analyzer Option for VSE/Power	CA MICS Analyzer Option for VSE/Power
NeuMICS Application Extension for Tandem	CA MICS Resource Management Tandem Option	CA MICS Tandem Option
NeuMICS MVS Batch and Operations Analyzer	CA MICS Resource Management Batch and Operations Analyzer	CA MICS Batch and Operations Analyzer
NeuMICS Application Extension for ASTEX	CA MICS Resource Management CA ASTEX Option	CA MICS CA ASTEX Option
NeuMICS Capacity Planner	CA MICS Resource Management Capacity Planner Option	CA MICS Capacity Planner Option
VM Data Transfer Program	CA MICS Resource Management Data Transfer Option for VM/CMS	CA MICS Data Transfer Option for VM/CMS
NeuMICS MVS Hardware and SCP Analyzer	CA MICS Resource Management Hardware and SCP Analyzer	CA MICS Hardware and SCP Analyzer
IMS Log Extractor	CA MICS Resource Management IMS Log Data Extractor Option	CA MICS IMS Log Data Extractor Option

PRODUCT'S OLD NAME	PRODUCT'S NEW FULL NAME	PRODUCT'S NEW SHORT NAME
NeuMICS SNA Network Analyzer	CA MICS Resource Management Network Analyzer Option	CA MICS Network Analyzer Option
NeuMICS MVS Performance Manager	CA MICS Resource Management Performance Manager Option	CA MICS Performance Manager Option
NeuMICS Q&R WorkStation	CA MICS Resource Management Q&R Workstation	CA MICS Q&R Workstation
NeuMICS DASD Space Analyzer	CA MICS Resource Management Space Analyzer Option	CA MICS Space Analyzer Option
NeuMICS DASD Space Collector	CA MICS Resource Management Space Collector Option	CA MICS Space Collector Option
NeuMICS StorageMate	CA MICS Resource Management StorageMate Option	CA MICS StorageMate Option
NeuMICS System Reliability Analyzer	CA MICS Resource Management System Reliability Analyzer Option	CA MICS System Reliability Analyzer Option
NeuMICS Web Analyzer	CA MICS Resource Management Web Analyzer Option	CA MICS Web Analyzer Option
N/A	CA MICS Resource Management Tape Analyzer Option	CA MICS Tape Analyzer Option

G E N E R A L T A B L E O F C O N T E N T S

	PAGE
Chapter 1. CA MICS ANALYZER FOR MEASUREWARE	1-001
1.1 Primary Areas of Application	1-001
1.2 Reporting and Inquiry Facilities	1-002
1.3 Files Overview	1-002
1.4 Product Prerequisites	1-003
Chapter 2. ANALYZER USAGE GUIDELINES	2-001
2.1 Data Analysis	2-001
2.2 CA MICS Product Interfaces	2-022
CHAPTER 3. MW OS AGENT ANALYZER REPORTS	3-001
3.1 MICF Inquiries for the MW OS Agent Analyzer	3-001
Chapter 4. RESERVED	4-001
Chapter 5. FILES	5-001
5.1 Data Element Naming Conventions	5-002
5.2 UNIX Global Performance File (AUMGPR)	5-002
5.3 UNIX Application Performance File (AUMAPP)	5-007
5.4 UNIX Process Activity File (AUMKPS)	5-010
5.5 UNIX Process Exception File (AUMIPS)	5-013
5.6 UNIX LAN Activity File (AUMNIF)	5-017
5.7 UNIX Disk Activity File (AUMDDA)	5-020
5.8 UNIX System and Collector Profile (AUMMCF)	5-023
5.9 UNIX Logical Volume Activity File (AUMVLM)	5-027
5.10 UNIX Business Transaction SLO Activity File (AUMTTA)	5-030
Chapter 6. DATA SOURCES	6-001
6.1 Data Collection Considerations	6-001
6.2 Record Descriptions	6-006
Chapter 7. PARAMETERS	7-001
7.1 Environmental Considerations	7-001
7.2 Complex Level Parameters	7-002
7.3 Unit Level Parameters	7-008
7.4 Other Related Parameters	7-043
CHAPTER 8. INSTALLATION	8-001
CHAPTER 9. MW OS AGENT ANALYZER PROCESSING	9-001
9.1 Processing Considerations	9-001
9.2 Daily Update Processing Flow	9-002
9.3 MICSLOG Excerpt from DAY089	9-003
Chapter 10. MODIFICATION	10-001
10.1 Standard User Exits	10-001
Appendix A. MESSAGES	A-001
Appendix B. DATA DICTIONARY	B-001
B.1 Common Data Elements	B-001
B.2 The UNIX Global Performance (AUMGPR) File	B-023
B.3 The UNIX Application Performance (AUMAPP) File	B-065
B.4 The UNIX Process Activity (AUMKPS) File	B-088
B.5 The UNIX Process Exception (AUMIPS) File	B-096
B.6 The UNIX LAN Activity (AUMNIF) File	B-108
B.7 The UNIX Disk Activity (AUMDDA) File	B-113
B.8 The UNIX System and Collector Profile (AUMMCF) File	B-124
B.9 The UNIX Logical Volume Activity (AUMVLM) File	B-136
B.10 The UNIX Bus. Transaction SLO Activity (AUMTTA) File	B-141

DETAIL TABLE OF CONTENTS PAGE

Chapter 1. CA MICS ANALYZER FOR MEASUREWARE 1-001

1.1 Primary Areas of Application 1-001

1.2 Reporting and Inquiry Facilities 1-002

1.3 Files Overview 1-002

1.4 Product Prerequisites. 1-003

Chapter 2. ANALYZER USAGE GUIDELINES 2-001

2.1 Data Analysis. 2-001
About MeasureWare Metrics 2-002
About the Analyzer for MeasureWare. 2-003
Analyzing UNIX Measurements 2-004

2.2 CA MICS Product Interfaces 2-022

CHAPTER 3. MW OS AGENT ANALYZER REPORTS. 3-001

3.1 MICF Inquiries for the MW OS Agent Analyzer. . . 3-001
3.1.1 UNIX Multiple System Inquiries 3-003
3.1.1.1 UNIX Global CPU Utilization (AUMCS1) 3-003
3.1.1.2 UNIX Global Queue Depth Comparison (AUMCS2). 3-005
3.1.2 UNIX Single System Inquiries 3-007
3.1.2.1 UNIX Global Bottlenecks (AUMCS4) 3-008
3.1.2.2 UNIX Global Disk Summary (AUMCS6). 3-010
3.1.3 UNIX Application Inquiries 3-012
3.1.3.1 UNIX Application Performance Summary (AUMLS2) 3-013
3.1.4 UNIX Process Activity Analysis Inquiries . . . 3-017
3.1.4.1 UNIX Detail Killed Process (AUMLS3). 3-017
3.1.5 UNIX Disk Inquiries. 3-022
3.1.5.1 UNIX Top 5 Disk Utilization Comparison (AUMCSB) 3-022
3.1.5.2 UNIX Disk I/O Activity (AUMCSC). 3-024
3.1.5.3 UNIX Logical Volume Activity (AUMCSD). . . . 3-026
3.1.6 UNIX Global Inquiries. 3-028
3.1.6.l UNIX System Performance Summary (AUMLS1) . . 3-029

Chapter 4. RESERVED. 4-001

Chapter 5. FILES 5-001

5.1 Data Element Naming Conventions. 5-002

5.2 UNIX Global Performance File (AUMGPR). 5-002
5.2.1 AUMGPR File Organization 5-003
5.2.2 AUMGPR Data Elements List. 5-003
5.2.3 AUMGPR Usage Considerations. 5-006
5.2.4 AUMGPR Retrieval Examples. 5-006

5.3 UNIX Application Performance File (AUMAPP) . . . 5-007
5.3.1 AUMAPP File Organization 5-007
5.3.2 AUMAPP Data Elements List. 5-008
5.3.3 AUMAPP Usage Considerations. 5-009
5.3.4 AUMAPP Retrieval Examples. 5-010

5.4 UNIX Process Activity File (AUMKPS). 5-010
5.4.1 AUMKPS File Organization 5-011
5.4.2 AUMKPS Data Elements List. 5-011
5.4.3 AUMKPS Usage Considerations. 5-012
5.4.4 AUMKPS Retrieval Examples. 5-013

5.5 UNIX Process Exception File (AUMIPS) 5-013
5.5.1 AUMIPS File Organization 5-014
5.5.2 AUMIPS Data Elements List. 5-014
5.5.3 AUMIPS Usage Considerations. 5-016
5.5.4 AUMIPS Retrieval Examples. 5-016

5.6 UNIX LAN Activity File (AUMNIF). 5-017
5.6.1 AUMNIF File Organization 5-017
5.6.2 AUMNIF Data Elements List. 5-018
5.6.3 AUMNIF Usage Considerations. 5-019
5.6.4 AUMNIF Retrieval Examples. 5-019

5.7 UNIX Disk Activity File (AUMDDA) 5-020
5.7.1 AUMDDA File Organization 5-020
5.7.2 AUMDDA Data Elements List. 5-021
5.7.3 AUMDDA Usage Considerations. 5-022
5.7.4 AUMDDA Retrieval Examples. 5-022

5.8 UNIX System and Collector Profile (AUMMCF) . . . 5-023
5.8.1 AUMMCF File Organization 5-024
5.8.2 AUMMCF Data Elements List. 5-024
5.8.3 AUMMCF Usage Considerations. 5-026
5.8.4 AUMMCF Retrieval Examples. 5-026

5.9 UNIX Logical Volume Activity File (AUMVLM) . . . 5-027
5.9.1 AUMVLM File Organization 5-027
5.9.2 AUMVLM Data Elements List. 5-028
5.9.3 AUMVLM Usage Considerations. 5-029
5.9.4 AUMVLM Retrieval Examples. 5-029

5.10 UNIX Business Transaction SLO Activity File (AUMTTA). 5-030
5.10.1 AUMTTA File Organization. 5-030
5.10.2 AUMTTA Data Elements List 5-031

DETAIL TABLE OF CONTENTS	PAGE
5.10.3 AUMTTA Usage Considerations	5-032
5.10.4 AUMTTA Retrieval Examples	5-033
Chapter 6. DATA SOURCES	6-001
6.1 Data Collection Considerations	6-001
6.2 Record Descriptions	6-006
Chapter 7. PARAMETERS	7-001
7.1 Environmental Considerations	7-001
7.2 Complex Level Parameters	7-002
7.2.1 Analyzer Definition Statements (AUMGENIN)	7-002
7.2.2 Account Code Structure (AUMACCT)	7-003
7.2.3 Account Code Derivation Routine (AUMACRT)	7-006
7.3 Unit Level Parameters	7-008
7.3.1 Processing Options (AUMOPS)	7-008
7.3.1.1 OPTIONS Statement	7-009
7.3.1.2 PLATFORM Statement	7-010
7.3.1.3 COMMON Statement	7-011
7.3.1.4 GLOBAL FORCE Statement	7-012
7.3.1.5 VERIFY Statement	7-012
7.3.1.6 WORK, MULTWORK, and NOMULT Statements	7-013
7.3.1.6.1 Change the Number of Work Files	7-017
7.3.1.7 Internal Step Restart Statements	7-018
7.3.1.7.1 Enable Internal Step Restart	7-022
7.3.1.8 Incremental Update Statements	7-023
7.3.1.8.1 Implement Incremental Update	7-033
7.3.1.9 Sample AUMOPS and Examples	7-034
7.3.2 Input DD Statements (INPUTAUM)	7-037
7.3.3 Input Needed to Generate System Code (AUMPGEN)	7-039
7.3.4 Application Name Derivation Routine (AUMAURT)	7-039
7.3.5 Database Space Modeling (DBMODEL)	7-041
7.3.5.1 Data Retention Specifications (FILE Statements)	7-041
7.3.5.2 DBMODEL Input Statements	7-043
7.4 Other Related Parameters	7-043
CHAPTER 8. INSTALLATION	8-001
CHAPTER 9. MW OS AGENT ANALYZER PROCESSING	9-001
9.1 Processing Considerations	9-001
9.2 Daily Update Processing Flow	9-002

DETAIL TABLE OF CONTENTS	PAGE
9.3 MICSLOG Excerpt from DAY089	9-003
Chapter 10. MODIFICATION	10-001
10.1 Standard User Exits	10-001
10.1.1 Input Exits	10-003
10.1.2 Parameter-Related Exits	10-004
10.1.3 Output Exits	10-005
Appendix A. MESSAGES	A-001
Appendix B. DATA DICTIONARY	B-001
B.1 Common Data Elements	B-001
AUMACTX Account Code Structure	B-001
AUMAPNUM Application Number	B-002
AUMAPU Application Name	B-002
AUMCPMDL System Model Identification	B-002
AUMGRPID Group Identification Number	B-002
AUMLDEV Logical Device Number	B-003
AUMLGVER Measurement Interface Version Number	B-003
AUMLNTTY Logon TTY Name	B-003
AUMOPSLC Operating System License	B-003
AUMOPSVR Operating System Release Number	B-004
AUMOPSYS Operating System Name	B-004
AUMPIN Process Identification Number	B-004
AUMPPID Parent Identification Number	B-004
AUMPRI Dispatch Priority	B-005
AUMPROG Program Name	B-005
AUMSAMPL Number of Data Samples	B-005
AUMSCOPE Scope Collector Name and Version Number	B-005
AUMSYSID Name of Host System	B-006
AUMUSER User Logon Name	B-006
DAY Day of Month	B-006
DAYNAME Name of Day of Week	B-006
DURATION Recording Interval Time	B-007
ENDTS End Time Stamp	B-010
HOUR Hour of Day	B-014
INTERVLS Number of Recording Intervals	B-014
MONTH Month of Year	B-015
ORGSYSID Originating System Identification	B-015
STARTTS Start Time Stamp	B-016
SYSID System Identifier	B-019
WEEK Week of Year	B-020
YEAR Year of Century	B-020
ZONE Time Zone	B-021
B.2 The UNIX Global Performance (AUMGPR) File	B-023
GPRACCPU CPUs Online per System	B-026
GPRACTVM Active Virtual Memory Size	B-026

DETAIL	TABLE OF CONTENTS	PAGE
GPRAPRCT	Sum of Avg Processes Active.	B-026
GPRAVAPR	Avg Number of Processes Active	B-026
GPRAVBIQ	Avg Processes Blocked on Local Disks . .	B-027
GPRAVCPQ	Avg Processes in Run Queue	B-027
GPRAVDKQ	Avg Processes Waiting for Disk I/Os. . .	B-028
GPRAVIPQ	Avg Processes Waiting for IPC Resources.	B-028
GPRAVLCL	Rate of LAN Collisions per Minute. . . .	B-028
GPRAVLER	Rate of LAN Errors per Minute.	B-029
GPRAVMMQ	Avg Processes Waiting for Memory	B-029
GPRAVNTQ	Avg Processes Waiting for Network. . . .	B-029
GPRAVPRO	Avg Number of Processes Alive.	B-030
GPRAVPRQ	Avg Processes Blocked on Priority. . . .	B-030
GPRAVRTM	Avg Run Time of Completing Processes . .	B-030
GPRAVSLQ	Avg Processes in Sleep State	B-030
GPRAVSUS	Avg Number of Processes Suspended. . . .	B-031
GPRAVSWA	Avg Number of Processes Waiting Swap-In.	B-031
GPRAVUSR	Avg Number of Login Users.	B-031
GPRBIOCT	Processes Blocked on Local Disks	B-032
GPRCACHE	Size of Buffer Cache Configured.	B-032
GPRCHUSE	The Size of Currently Used Buffers . . .	B-032
GPRCPQTM	CPU Queue Wait Time.	B-032
GPRCPRCT	Total Processes Completed.	B-033
GPRCPUCP	Total Available CPU Time	B-033
GPRCPUTM	Total CPU Busy	B-033
GPRCSWTM	CPU Busy for Context Switching	B-033
GPRDBKCT	Block I/Os	B-034
GPRDBRCT	Block Reads.	B-034
GPRDBWCT	Block Writes	B-034
GPRDFRCT	Disk Reads by File System.	B-035
GPRDFWCT	Disk Writes by File System	B-035
GPRDKLRT	Logical Disk I/O Rate.	B-035
GPRDKQTM	Disk I/O Queue Wait Time	B-035
GPRDMMCT	Virtual Memory Management I/Os	B-036
GPRDMMRT	Virtual Memory Management I/O Rate . . .	B-036
GPRDMRCT	Disk Reads for Virt Memory Management. .	B-036
GPRDMRRT	Virt Memory Management Disk Read Rate. .	B-037
GPRDMWCT	Disk Writes for Virt Memory Management .	B-037
GPRDMWRT	Virt Memory Management Disk Write Rate .	B-037
GPRDRRCT	Raw Reads.	B-038
GPRDRWCT	Raw Writes	B-038
GPRDSKIO	Physical Disk I/Os	B-038
GPRDSKRT	Physical Disk I/O Rate	B-039
GPRDSRCT	Disk Reads for File System Management. .	B-039
GPRDSWCT	Disk Writes for File System Management .	B-039
GPRFSDCT	File System I/Os	B-039
GPRFSMCT	File System Management I/Os.	B-040
GPRINODU	Inode Entries in Use	B-040
GPRINTTM	CPU Busy for Interrupt Handling.	B-040
GPRIPQTM	IPC Queue Wait Time.	B-040
GPRLCOLS	LAN Collisions	B-041

DETAIL	TABLE OF CONTENTS	PAGE
GPRLDKCT	Logical Disk I/Os.	B-041
GPRLERRS	LAN Packet Errors.	B-041
GPRLOSTB	Trace Buffers Lost	B-041
GPRLPCKI	Successful LAN Packets Received.	B-042
GPRLPCKO	Successful LAN Packets Sent.	B-042
GPRLRDCT	Logical Disk Reads	B-042
GPRLRDKB	KBytes Logically Read From Disk.	B-043
GPRLUSED	Number of File Locks in Use.	B-043
GPRLWTCT	Logical Disk Writes.	B-043
GPRLWTKB	KBytes Logically Written To Disk	B-043
GPRMEMSW	Amount of Swap Space Configured.	B-044
GPRMEMSZ	Amount of Main Memory Configured	B-044
GPRMEMUS	Amount of Main Memory Available for User	B-044
GPRMEMUT	Main Memory Used	B-044
GPRMMEMS	Main Memory Used by System	B-045
GPRMMEMU	Main Memory Used by Users.	B-045
GPRMMQTM	Memory Queue Wait Time	B-045
GPRMPBDU	Percent Utilization of Busiest Disk. . .	B-045
GPRMPDSU	Peak Disk Space Utilization.	B-046
GPRMQUSE	Message Queues in Use.	B-046
GPRMSGQU	Message Queues Available	B-046
GPRMXCPU	Max Number of Processors Configured. . .	B-047
GPRMXDSK	Max Number of Disk Drives Configured . .	B-047
GPRMXDTM	Time Busiest Disk Was Busy	B-047
GPRMXLAN	Max Number of LAN Interfaces Configured.	B-047
GPRNFSCT	Network File System Requests	B-048
GPRNFSRT	Rate of Network File System Requests . .	B-048
GPRNICTM	CPU Busy at Nice Priorities.	B-048
GPRNORTM	CPU Busy at Normal Priorities.	B-049
GPRNPROC	Value of NPROC Parameter	B-049
GPRNTQTM	Network Queue Wait Time.	B-049
GPROFLOW	Unmeasurable New Transactions.	B-049
GPROUTQU	Outbound Queue Lengths	B-050
GPRPAGRT	Page Fault Rate.	B-050
GPRPAUTM	CPU Idle Due to I/O Wait	B-050
GPRPCAVM	Percent of Virtual Memory in Active Use.	B-050
GPRPCCPU	Percent Total CPU Busy	B-051
GPRPCCSW	Percent CPU Busy for Context Switching .	B-051
GPRPCHIT	Cache Hit Ratio.	B-051
GPRPCINT	Percent CPU Busy for Interrupt Handling.	B-052
GPRPCMEM	Percent of Main Memory Used.	B-052
GPRPCMMS	Percent of Main Memory Used by System. .	B-052
GPRPCMMU	Percent of Main Memory Used by Users . .	B-053
GPRPCNIC	Percent CPU Busy at Nice Priorities. . .	B-053
GPRPCNOR	Percent CPU Busy at Normal Priorities. .	B-053
GPRPCRTM	Percent CPU Busy at Real Time Priorities	B-054
GPRPCSSC	Percent CPU Busy for System Calls. . . .	B-054
GPRPCSWS	Percent of Swap Space Used	B-054
GPRPCSYS	Percent CPU Busy in System Mode.	B-055
GPRPCUSR	Percent CPU Busy in User Mode.	B-055

DETAIL TABLE OF CONTENTS		PAGE
GPRPFLCT	Page Faults.	B-055
GPRPGOCT	Number of Page Outs.	B-055
GPRPKBCT	KBytes Physically Transferred.	B-056
GPRPRDCT	Physical Disk Reads.	B-056
GPRPRITM	Time Blocked on Priority	B-056
GPRPROCT	Sum of Avg Processes Alive	B-056
GPRPTBLU	Proc Table Entries Used.	B-057
GPRPWTCT	Physical Disk Writes	B-057
GPRRAWCT	Raw I/Os	B-057
GPRRDHIT	Buffered Read Hits	B-058
GPRRTMTM	CPU Busy at Real Time Priorities	B-058
GPRRUNTM	Total Run Time of Completing Processes .	B-058
GPRSCALL	Total System Calls	B-058
GPRSCLRT	System Procedure Call Rate	B-059
GPRSEMID	Semaphore Identifiers Available.	B-059
GPRSEMST	Semaphore Identifiers in Use	B-059
GPRSHMSG	Table Memory Segments Available.	B-060
GPRSLQTM	Sleep Queue Wait Time.	B-060
GPRSMEMS	Shared Memory Segments in Use.	B-060
GPRSPRCT	Total Processes Started.	B-060
GPRSSCTM	CPU Busy for System Calls.	B-061
GPRSUPTM	Time of Last System Reboot	B-061
GPRSUSCT	Procs Suspended During Interval.	B-061
GPRSWACT	Processes Waiting to be Swapped in . . .	B-061
GPRSWAPU	Swap Space Used.	B-062
GPRSWOCT	Number of Swap Outs/Deactivations. . . .	B-062
GPRSWORT	Swap-Out Rate.	B-062
GPRSWPCT	Process Swaps, Activation/Deactivation .	B-063
GPRSWPRT	Swapping Rate.	B-063
GPRSYSTM	CPU Busy in System Mode.	B-063
GPRTBLEN	File Table Entries Used.	B-063
GPRTBLUK	Entries Used by Kernel	B-064
GPRUSRCT	Sum of Avg Login Users	B-064
GPRUSRTM	CPU Busy in User Mode.	B-064
B.3	The UNIX Application Performance (AUMAPP) File .	B-065
APPAPRCT	Sum of Avg Processes Active.	B-067
APPAVAPR	Avg Processes Active	B-067
APPAVPRO	Avg Processes Alive.	B-068
APPAVRTM	Avg Run Time of Completing Processes . .	B-068
APPAVSUS	Avg Number of Processes Suspended. . . .	B-068
APPCOST	Processing Charges	B-068
APPCPRCT	Number of Processes Completed.	B-069
APPCPUCM	PRM CPU Cap Mode	B-069
APPCPUCP	Total Available CPU Time	B-069
APPCPUEN	PRM CPU Entitlement.	B-070
APPCPUST	PRM CPU State.	B-070
APPCPUTM	Total CPU Use.	B-070
APPDBKCT	Block I/Os	B-070
APPDBRCT	Block Reads.	B-071

DETAIL TABLE OF CONTENTS		PAGE
APPDBWCT	Block Writes	B-071
APPDFSCT	Number of File System Reads and Writes .	B-071
APPDFSRT	Rate of File System Reads/Writes (IO/Sec	B-072
APPDMMCT	Number of Memory Manager Reads and Write	B-072
APPDMMRT	Rate of Memory Mgr. Reads/Writes (IO/Sec	B-072
APPDRAIO	Number of Raw Reads and Writes	B-072
APPDRART	Rate of Raw Reads/Writes (IO/Sec). . . .	B-073
APPDSKIO	Number of Physical Disk I/Os	B-073
APPDSYIO	Number of System Reads and Writes. . . .	B-073
APPDSYRT	Rate of System Reads/Writes (IO/Sec) . .	B-073
APPLIOCT	Logical Disk I/Os.	B-074
APPLOGMD	PRM Logging Mode	B-074
APPLRDCT	Logical Disk Reads	B-074
APPLWTCT	Logical Disk Writes.	B-075
APPMAJRT	Rate of Page Faults to Disk.	B-075
APPMEMAV	PRM Available Memory	B-075
APPMEMEN	PRM MEM Entitlement.	B-075
APPMEMPR	Private Resident Memory Size in KBytes .	B-076
APPMEMRS	Resident Memory Size in KBytes	B-076
APPMEMST	PRM MEM State.	B-076
APPMEMSW	Amount of Swap Space Configured.	B-077
APPMEMSZ	Amount of Main Memory Configured	B-077
APPMEMUP	PRM MEM Upperbound	B-077
APPMEMUS	PRM Memory Used by Processes	B-077
APPMEMVR	Virtual Memory Size in KBytes.	B-078
APPMINRT	Rate of Page Faults to Memory.	B-078
APPNICTM	CPU Use at Nice Priorities	B-078
APPNORTM	CPU Use at Normal Priorities	B-078
APPPCCPU	Percent Total CPU Use.	B-079
APPPCMPR	Percent Memory Use for Private Resident.	B-079
APPPCMRS	Percent Memory Use for Resident Memory .	B-079
APPPCNIC	Percent CPU Use at Nice Priorities . . .	B-080
APPPCNOR	Percent CPU Use at Normal Priorities . .	B-080
APPPCRTM	Percent CPU Use at Real Time Priorities.	B-080
APPPCSCP	Pct CPU Wait	B-080
APPPCSDK	Pct Disk I/O Wait.	B-081
APPPCSDR	Pct Time Directed to Wait.	B-081
APPPCSIM	Pct Software Impede Wait	B-081
APPPCSIP	Pct Interprocess Communication Wait. . .	B-081
APPPCSLN	Pct LAN I/O Wait	B-082
APPPCSOT	Pct Non-LAN/DISK/TERM I/O Wait	B-082
APPPCSWP	Pct Memory Wait.	B-082
APPPCSYS	Percent CPU Use for System Processing. .	B-082
APPPCUSR	Percent CPU Use in User Mode	B-083
APPPFMAJ	Page Faults to Disk.	B-083
APPPFMIN	Page Faults to Memory.	B-083
APPPKBCT	I/O Transfer Count in KBytes	B-083
APPPKBRT	I/O Transfer Rate in KBytes.	B-084
APPPRDCT	Physical Disk Reads.	B-084
APPPROCT	Sum of Avg Processes Alive	B-084

DETAIL TABLE OF CONTENTS		PAGE
APPPSPRD	Physical Disk Read Rate.	B-085
APPPSPWT	Physical Disk Write Rate	B-085
APPPWTCT	Physical Disk Writes	B-085
APPRTMTM	CPU Use at Real Time Priorities.	B-085
APPRUNTM	Total Run Time of Completing Processes .	B-086
APPSUSCT	Procs suspended during interval.	B-086
APPSYSTM	CPU Use for System Processing.	B-086
APPUSRTM	CPU Use in User Mode	B-086
B.4 The UNIX Process Activity (AUMKPS) File.		B-088
KPSAVDRT	Physical Disk I/O Rate	B-090
KPSAVLRT	Logical Disk I/O Rate.	B-090
KPSCOST	Processing Charges	B-090
KPSCPUTM	Total CPU Use.	B-090
KPSDBKCT	Block I/Os	B-091
KPSDBKRT	Block I/O Rate	B-091
KPSDSKIO	Physical Disk I/Os	B-091
KPSLDKCT	Logical Disk I/Os.	B-091
KPSPKBCT	Total I/O Transfer Count in KBytes . . .	B-092
KPSPKBRT	I/O Transfer Rate in KBytes.	B-092
KPSRSSKB	Resident Set Size.	B-092
KPSRUNTM	Total Run Time	B-092
KPSSCPTM	CPU Wait Time.	B-093
KPSSDKTM	Disk I/O Wait Time	B-093
KPSSDRTM	Directed to Wait Time.	B-093
KPSSIMTM	Software Impede Wait Time.	B-093
KPSSIPTM	IPC Resource Wait Time	B-094
KPSSLNTM	LAN I/O Wait Time.	B-094
KPSSNFTM	Network File System Wait Time.	B-094
KPSSOITM	Non-LAN/DISK/TERM I/O Wait Time.	B-094
KPSSSYTM	System Wait Time	B-095
KPSSVMTM	Virtual Memory Wait Time	B-095
KPSVSSKB	Virtual Memory Regions	B-095
B.5 The UNIX Process Exception (AUMIPS) File		B-096
IPSALITM	Process Alive Time	B-098
IPSCPUTM	Total CPU Use.	B-098
IPSCSWTM	CPU Use for Context Switching.	B-098
IPSCUMTM	Cumulative To-Date CPU Time.	B-098
IPSDBKCT	Block I/Os	B-099
IPSDBRCT	Block Reads.	B-099
IPSDBWCT	Block Writes	B-100
IPSDFRCT	File System Disk Reads	B-100
IPSDFWCT	File System Disk Writes.	B-100
IPSDMMCT	Virtual Memory I/Os.	B-100
IPSDSKIO	Physical Disk I/Os	B-101
IPSDSYCT	File System Management I/Os.	B-101
IPSFSDCT	File System Disk I/Os.	B-101
IPSINTR	Interest Reason Codes.	B-101
IPSINTTM	CPU Use for Interrupt Handling	B-102

DETAIL TABLE OF CONTENTS		PAGE
IPSLRDCT	Logical Disk Reads	B-102
IPSLWTCT	Logical Disk Writes.	B-102
IPSNICTM	CPU Use at Nice Priority	B-103
IPSNORTM	CPU Use at Normal Priority	B-103
IPSPFMAJ	Page Faults to Disk.	B-103
IPSPFMIN	Page Faults to Memory.	B-103
IPSPKBCT	I/O Transfer Count in KBytes	B-104
IPSPRMID	PRM Group Identification Number.	B-104
IPSRSSKB	Resident Set Size.	B-104
IPSRTMTM	CPU Use at Real Time Priority.	B-104
IPSRUNTM	Total Accumulated Run Time	B-105
IPSSSCTM	CPU Use for System Calls	B-105
IPSSTOPR	Blocked State.	B-105
IPSSYSTM	CPU Use in System Mode	B-106
IPSTHCNT	Total Number of Thread Counts.	B-106
IPSUSRTM	CPU Use in User Mode	B-106
IPSVSSKB	Virtual Memory Regions	B-107
B.6 The UNIX LAN Activity (AUMNIF) File.		B-108
NIFCOLCT	LAN Collisions	B-109
NIFERRCT	LAN Errors	B-110
NIFMTUSZ	MTU Size	B-110
NIFNAME	Device Name for LAN Interface.	B-110
NIFPKICT	Inbound LAN Packets.	B-110
NIFPKOCT	Outbound LAN Packets	B-110
NIFPSCOL	LAN Collision Rate	B-111
NIFPSERR	LAN Error Rate	B-111
NIFPSPKI	Inbound LAN Packet Rate.	B-111
NIFPSPKO	Outbound LAN Packet Rate	B-111
NIFQUEUE	Outbound Queue Length.	B-112
NIFSPEED	Interface Speed.	B-112
B.7 The UNIX Disk Activity (AUMDDA) File		B-113
DDAAVQLN	Avg Queue Length	B-115
DDADDPDR	Time Processing Disk Request	B-115
DDADFRCT	File System Reads.	B-115
DDADFWCT	File System Writes	B-115
DDADIR	File System Directory.	B-116
DDADMMCT	Virtual Memory I/Os.	B-116
DDADRRCT	Raw Reads.	B-116
DDADRWCT	Raw Writes	B-116
DDADSKIO	Physical I/Os.	B-116
DDADSKRT	Physical I/O Rate.	B-117
DDADUTTM	Disk Busy Time	B-117
DDALRDCT	Logical Disk Reads	B-117
DDALWTCT	Logical Disk Writes.	B-118
DDANAME	Disk Name.	B-118
DDAPCDUT	Percent Disk Utilization	B-118
DDAPKBCT	Disk I/O Transfer Count in KBytes. . . .	B-118
DDAPKBRT	Disk I/O Transfer Rate in KBytes	B-119

DDAPRDCT	Physical Reads	B-119
DDAPSFRD	File System Read Rate	B-119
DDAPSFWT	File System Write Rate	B-119
DDAPSLRD	Logical Disk Read Rate	B-120
DDAPSLWT	Logical Disk Write Rate	B-120
DDAPSMEM	Virtual Memory I/O Rate	B-120
DDAPSRRD	Raw Read Rate	B-121
DDAPSRWT	Raw Write Rate	B-121
DDAPSSYS	File System Management I/O Rate	B-121
DDAPWTCT	Physical Writes	B-121
DDAQLNSM	Sum of Avg Queue Length	B-122
DDARKBCT	KBytes Read	B-122
DDARKBRT	Disk Read Transfer Rate in KBytes	B-122
DDASYSCT	File System Management I/Os	B-123
DDAWKBCT	KBytes Written	B-123
DDAWKBRT	Disk Write Transfer Rate in KBytes	B-123
B.8	The UNIX System and Collector Profile (AUMMCF) File	B-124
MCFCACHE	Size of Buffer Cache Configured	B-125
MCFCPUMN	Threshold, Percent of CPU	B-125
MCFINODE	Entries in Inode Table on System	B-126
MCFIORMN	Threshold, Physical Disk I/O Rate	B-126
MCFKRNSZ	Kernel Word Size	B-126
MCFLCKAV	File/Record Locks Available	B-126
MCFLGTPS	Data Types Logged by Collector	B-127
MCFLOSHR	Flag - Log Short-Running Processes	B-127
MCFMCMOD	CPU Model	B-127
MCFMEMSW	Swap Space Configured	B-127
MCFMEMSZ	Main Memory Configured	B-128
MCFMEMUS	Main Memory Available for Users	B-128
MCFMNCHC	Minimum Size of Buffer Cache Configured	B-128
MCFMNCPU	Minimum Number of Processors Configured	B-128
MCFMNDSK	Minimum Number of Disk Drives Configured	B-128
MCFMNLAN	Minimum Num of LAN Interfaces Configured	B-129
MCFMNMEM	Minimum Main Memory Configured	B-129
MCFMNMUS	Minimum Main Memory Available for Users	B-129
MCFMNSWS	Minimum Swap Space Configured	B-129
MCFMSGQU	Message Queues Available	B-130
MCFMXCHC	Maximum Size of Buffer Cache Configured	B-130
MCFMXCPU	Maximum Number of Processors Configured	B-130
MCFMXDSK	Maximum Number of Disk Drives Configured	B-130
MCFMXLAN	Maximum Num of LAN Interfaces Configured	B-130
MCFMXMEM	Maximum Main Memory Configured	B-131
MCFMXMUS	Maximum Main Memory Available for Users	B-131
MCFMXSWS	Maximum Swap Space Configured	B-131
MCFNCPUS	Number of Processors Configured	B-131
MCFNDSKS	Number of Disk Drives Configured	B-131
MCFNLANS	Number of LAN Interfaces Configured	B-132
MCFNOKIL	Flag - No Terminating Processes	B-132
MCFNONEW	Flag - No Newly Created Processes	B-132
MCFNPROC	Value of NPROC Parameter	B-132
MCFOPSYS	Operating System Name	B-133
MCFPRCMM	Threshold, Virtual Memory	B-133
MCFSEMID	Semaphore Identifiers Available	B-133
MCFSHMSG	Table Memory Segments Available	B-133
MCFTBLUK	Entries Used by Kernel	B-134
MCFWCPMN	Threshold, Percent of CPU Wait	B-134
MCFWDKMN	Threshold, Percent of Disk Wait	B-134
MCFWIMMN	Threshold, Percent of Semaphore Wait	B-134
MCFWMEMN	Threshold, Percent of Memory Wait	B-135
B.9	The UNIX Logical Volume Activity (AUMVLM) File	B-136
VLMDIR	Logical Volume Directory	B-137
VLMLRDCT	Logical Disk Reads	B-138
VLMLWTCT	Logical Disk Writes	B-138
VLMMNSPU	Minimum Pct Disk Space Used	B-138
VLMMXSPU	Maximum Pct Disk Space Used	B-138
VLMNAME	Logical Volume Group Name	B-139
VLMPCSPU	Percent Disk Space Used	B-139
VLMPSLRD	Logical Disk Read Rate	B-139
VLMPSLWT	Logical Disk Write Rate	B-139
VLMPTHNM	Mount Point Path Name	B-140
VLMSPUSM	Sum of Avg Disk Space Used	B-140
B.10	The UNIX Bus. Transaction SLO Activity (AUMTTA) File	B-141
TTAABORT	Aborted Transaction Count	B-142
TTAABTTM	Cumulative Aborted Transaction Time	B-142
TTAADFMT	Correlator Address Format	B-143
TTAARMID	ARM Transaction ID	B-143
TTAARMIF	ARM Transaction Information	B-143
TTAARMNM	ARM Application Name	B-143
TTAARMUN	ARM Transaction User Name	B-144
TTAAVATM	Average Time Per Aborted Trans	B-144
TTAAVTTM	Average Transaction Response Time	B-144
TTABINCT	Number of Transaction Bins Defined	B-145
TTACORAD	Correlator Address	B-145
TTACPUTM	Total CPU Consumed by Transaction	B-145
TTADSKIO	Physical Disk I/Os per Transaction	B-145
TTALGLIO	Logical I/Os per Transaction	B-145
TTANAME	Transaction Name	B-146
TTAPCBXX	Percent Transactions in Bin 1	B-146
TTAPCVSL	Percent Transactions Violated SLO	B-146
TTASLOMX	Transaction Service Level Objective	B-147
TTATFAIL	Number of Failed Transactions	B-147
TTATRANS	Completed Transaction Count	B-147
TTATRNID	Transaction Class ID	B-147
TTATRNTM	Cumulative Transaction Response Time	B-147
TTATRNXX	Completed Transactions in Bin 1	B-148

D E T A I L T A B L E O F C O N T E N T S PAGE

TTAUPRXX Upper Range Response Time for Bin 1. . . B-148
TTAVSLCT Violated SLO Transaction Count B-148

F I G U R E T A B L E O F C O N T E N T S

NUMBER	TITLE	PAGE
2-0	Common Metrics	2-005
2-1	Global System Data (Part 1 of 4)	2-006
2-1	Global System Data (Part 2 of 4)	2-007
2-1	Global System Data (Part 3 of 4)	2-008
2-1	Global System Data (Part 4 of 4)	2-009
2-2	Application Data (Part 1 of 3)	2-010
2-2	Application Data (Part 2 of 3)	2-011
2-2	Application Data (Part 3 of 3)	2-012
2-3	Killed Process Data (Part 1 of 2).	2-013
2-3	Killed Process Data (Part 2 of 2).	2-014
2-4	Process Exception Data	2-015
2-5	LAN Activity Data.	2-016
2-6	Disk Activity Data	2-017
2-7	Logical Volume Data.	2-018
2-8	Business Transaction SLO Activity Data . .	2-019
2-9	System Configuration & Collector Profile Data (Part 1 of 2).	2-020
2-9	System Configuration & Collector Profile Data (Part 2 of 2).	2-021
5-1	UNIX Information Area Files.	5-001
5-2	AUMGPR Timespan Granularity Chart.	5-003
5-3	AUMAPP Timespan Granularity Chart.	5-007
5-4	AUMKPS Timespan Granularity Chart.	5-011
5-5	AUMIPS Timespan Granularity Chart.	5-014
5-6	AUMNIF Timespan Granularity Chart.	5-017
5-7	AUMDDA Timespan Granularity Chart.	5-020
5-8	AUMMCF Timespan Granularity Chart.	5-024
5-9	AUMVLM Timespan Granularity Chart.	5-027
5-10	AUMTTA Timespan Granularity Chart.	5-030
6-1	Record Descriptions (Part 1 of 2).	6-006
6-1	Record Descriptions (Part 2 of 2).	6-007
7-1	Account Code Level Definition Worksheet. .	7-005
7-2	Account Code Derivation Routine Worksheet.	7-007
7-3	AUMOPS Options Definition Worksheet. . . .	7-036
7-4	INPUTAUM JCL Definitions Worksheet	7-038
7-5	Members/Macros Needed/Updated by AUMPGEN Code.	7-039
7-6	Application Name Derivation Routine Worksheet	7-040
7-7	Database Data Retention Definitions Worksheet	7-042
9-1	MICSLOG Extract from DAY089 for an HP-UX System at B.10 (Part 1 of 3).	9-004
9-1	MICSLOG Extract from DAY089 for an HP-UX System at B.10 (Part 2 of 3).	9-005
9-1	MICSLOG Extract from DAY089 for an HP-UX System at B.10 (Part 3 of 3).	9-006

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 1

OVERVIEW

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 1

OVERVIEW

Chapter 1. CA MICS ANALYZER FOR MEASUREWARE

As the number of management information systems installations using UNIX systems continues to grow, the need to systematically plan, control, and evaluate the environment is becoming increasingly important.

CA has responded to this need by developing the CA MICS Analyzer for MeasureWare (referred to as MW OS Agent Analyzer), a data integration application that is part of the CA MICS product suite. Like other data integration applications, the MW OS Agent Analyzer processes raw data for analyzing and managing specific technologies.

The MW OS Agent Analyzer uses data from Hewlett-Packard's MeasureWare (renamed by HP to OpenView) Agent's logglob, logappl, logproc, logdev, and logtran files to populate the CA MICS database's UNIX information area (AUM). Because this data is stored in the CA MICS database in a format common to all CA MICS products, you can use standard CA MICS facilities to integrate MW OS Agent Analyzer data with any other CA MICS management support applications such as financial and capacity management.

CA MICS offers online data access through the CA MICS Information Center Facility (MICF), which is a full-screen, menu-based processor that allows you to perform data analysis, inquiry, retrieval, and report functions.

The MW OS Agent Analyzer provides color graphics reports, using MICF, that enable you to combine data from multiple UNIX systems and to report on global system performance and workload, process activity, interesting processes, disks, and the profile of your system and its configuration.

This chapter discusses the MW OS Agent Analyzer in the following sections:

1 - Primary Areas of Application
2 - Reporting and Inquiry Facilities
3 - Files Overview
4 - Product Prerequisites

1.1 - Primary Areas of Application

The CA MICS Analyzer for MeasureWare (MW OS Agent Analyzer) provides data and reporting capabilities that aid in managing your UNIX environment in the following areas:

- Capacity Planning - The MW OS Agent Analyzer maintains a database of historical MeasureWare usage and performance data. This information can be used for growth and trend projections at both the system and organizational unit (such as division or department) levels.

- Performance Management - The MW OS Agent Analyzer enables you to examine and analyze specific workloads so you can optimize your resources. This information can be used for monitoring system behavior on a routine basis, investigating user reports of degraded performance, performing system tuning, and assessing the effects of changing system software and hardware configurations.

- Accounting and Chargeback - A standard interface to CA MICS Accounting and Chargeback enables you to use the Analyzer's UNIX Application Performance and UNIX Process Activity files to bill for UNIX system use. This data includes CPU time, I/O counts, and other resource usage information.

- Service Level Management - The MW OS Agent Analyzer enables you to use its data to create and maintain service level agreements (SLAs). Creating SLAs and reporting on the success of IS in meeting them requires that application data be in an easily accessible format. The Analyzer accomplishes this by combining data into user-specified groups (such as accounts payable) to provide consolidated reporting on SLA achievement.

- Management and Analytical Reporting - The MW OS Agent Analyzer enables you to group and summarize information directly from the CA MICS database to create useful ad hoc management and analytical reports. CA MICS facilities allow resummarization when the data in the CA MICS database does not address specific reporting needs.

1.2 - Reporting and Inquiry Facilities

The MW OS Agent Analyzer enables you to use the MICF and standard SAS language interfaces to create ad hoc reports. CA MICS provides these interactive capabilities, which enable you to quickly respond to complex requests, so that you have the information you need to manage your IS organization each day. These capabilities are provided by:

- MICF--MICF is a panel-oriented productivity tool that helps you to access information in the CA MICS database. When you define input, selection criteria, and report options, MICF dynamically builds an inquiry program that executes in either interactive or batch mode. Then, depending on your specific requirements, MICF either prints the inquiry report, displays the results at your terminal, or catalogs the results for later viewing.

- The Standard SAS Language Interface--The CA MICS Workstation Facility (MWF) provides an online environment that supports the standard SAS language interface in both interactive and batch modes. SAS allows easy data manipulation and statistical analysis and provides coding facilities for programmers who are conducting extensive analysis efforts or designing new reports for IS business applications.

1.3 - Files Overview

The MW OS Agent Analyzer builds the UNIX information area. Once created, this information area maintains system and user workload and utilization data.

File	Contents and Use
UNIX Global System Performance (AUMGPR)	Global system measurements; capacity, system load, and performance analyses
UNIX Application Performance (AUMAPP)	Interval-based measures; application use
UNIX Process Activity (AUMKPS)	Completed process measures; accounting and process analysis
UNIX Process Exception (AUMIPS)	Interesting process measures; performance and problem analyses; modeling
UNIX LAN Activity (AUMNIF)	Interval-based LAN measures; capacity analysis
UNIX Disk Activity (AUMDDA)	Global system measures; capacity analysis
UNIX System and Collector Profile (AUMMCF)	Global system measures; problem diagnosis, change tracking, reporting
UNIX Logical Volume Activity (AUMVLM)	Interval-based I/O and space measures; capacity analysis
UNIX Business Transaction SLO Activity (AUMTTA)	Interval-based business transaction measures; service level analysis

1.4 - Product Prerequisites

The CA MICS Analyzer for MeasureWare runs in a CA MICS environment that operates under z/OS.

The data for the Analyzer comes from HP MeasureWare for HP-UX Systems (HP-UX). The Analyzer supports:

Your Environment and Release	scopeux Release
HP-UX 9.0	B.02.02
HP-UX 10.0	B.02.02
HP-UX 10.10 - 10.30	C.00.01
HP-UX 10.20 & 10.30	C.01.00 through C.02.15
HP-UX 10.20 & 11.0	C.02.30 through C.02.60
HP-UX 10.20, 11.0, and 11i	C.03.20 through C.03.50
HP-UX 11.0 & 11.11	C.03.58 through C.03.86
AIX 3.2	B.02.04
AIX 4.1.x and 4.2.x	B.02.19
AIX 4.3.1, 4.3.2, and 4.3.3	C.02.58 through C.03.15
AIX 4.3.2, 4.3.3, and 5L V5.1	C.03.40
AIX 4.3.2, 4.3.3, 5L V5.1 and V5.2	C.03.80
AIX 4.3.2, 4.3.3, 5L V5.1, V5.2, and V5.3	C.04.5
Solaris 2.5 & 2.5.1 (SunOS 5.4 & 5.6)	B.02.04
Solaris 2.5, 2.5.1, and 2.6 SunOS 5.4 and up	C.02.05 through C.02.06
Solaris 2.6, 7, and 8 SunOS 5.6 and up	C.02.53 through C.03.45
Solaris 2.6, 7, 8, and 9 SunOS 5.6 and up	C.03.75 and C.03.82

You may also want to consider installing SAS/GRAPH because some of the MICF inquiries distributed with the Analyzer use its facilities.

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 2

USAGE GUIDELINES

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 2

USAGE GUIDELINES

Chapter 2. ANALYZER USAGE GUIDELINES

The CA MICS Analyzer for MeasureWare provides:

- o Repository management services for metrics from Hewlett-Packard's MeasureWare Agent Software, enabling you to manage your enterprise's UNIX system data using the same CA MICS techniques you use to manage MVS, VM, and VAX/VMS data

- o Reports, using the CA MICS Information Center Facility (MICF), that enable you to manage the financial, capacity, and performance aspects of your enterprise's UNIX systems

The sections that follow describe the Analyzer's usefulness in greater depth:

1 - Data Analysis
2 - CA MICS Product Interfaces

2.1 - Data Analysis

To help you better understand how to use the data provided by the Analyzer for MeasureWare, this section first presents information about the data available from UNIX systems. It then presents tables that help you choose CA MICS data elements for managing your systems.

About MeasureWare Metrics

Hewlett-Packard's MeasureWare Agent Software contains the following types of metrics:

- Global metrics report on system activity in 5-minute and user-defined summaries.

  The kinds of measures that comprise global metrics include:

  - CPU (usage at various priorities/modes)
  - disk (read and write rates; peak disk use)
  - networking (global LAN errors and collisions)
  - memory use, paging, and swapping
  - process queue length (CPU, memory, disk, network, IPC, and so on)
  - user and process workload
  - terminal transaction (rate and response time)
  - DCE server interface (completed operations, server errors, RPC response time)

- Application metrics capture data by user-defined applications, which are groups of processes.

  The kinds of measures that comprise application metrics include:

  - application identification (application number and name)
  - summary (CPU use during interval)
  - process count (numbers of processes)
  - terminal transaction (number of and rates of terminal transactions)
  - process wait (time spent in wait states)
  - CPU (usage for different reasons)
  - disk (read and write rates)
  - memory (virtual memory size, resident memory size)

- Process metrics capture data about each process (a process is similar to a transaction in MVS terminology).

  The kinds of measures that comprise process metrics include:

  - process identification (user name, interest codes)
  - process summary (runtime)
  - CPU (CPU use at various priorities/modes)
  - disk (read and write rates)
  - memory (page faults, resident set size, and virtual memory size)
  - terminal transaction (number of and rates of terminal transactions)
  - process wait (time spent in wait states)

- LAN metrics capture data about each network interface.

  The kind of measures that comprise LAN metrics include:

  - individual inbound and outbound packets, error and collision rates.

- Disk metrics capture data about each disk.

  The kinds of measures that comprise disk metrics include:

  - logical and physical disk I/Os, average queue length, disk utilization, disk space use

- Volume metrics capture data about each logical volume.

  The kinds of measures that comprise volume metrics include:

  - logical disk I/Os, disk space use

- Transaction metrics capture data about each business transaction tracked by HP Transaction Tracker.

  The kind of measures that comprise transaction metrics include:

  - defined parameters (transaction service-level objective, number of transaction bins, upper range response time for each bin)

  - transaction activity summary (the percent of transactions that violated service-level objectives, completed and aborted transaction counts, transaction counts in various response time bins)

- Configuration metrics identify the system and collector on which the data was captured and the response time distributions and process thresholds established, among other measures.

  The kinds of measures that comprise configuration metrics include:

  - Record identification (date, time, length of measurement interval)

  - System and collector identification (system name,

collector version, numbers of processors, disks, and LANs configured)

- Interesting process thresholds (thresholds coded in the UNIX system's parm file)

Note: HP MeasureWare supports HP's Data Source Integration (DSI) package. The CA MICS Analyzer for MeasureWare does not support DSI data.

About the Analyzer for MeasureWare

The CA MICS Analyzer for MeasureWare helps you manage your UNIX environment by:

- o Providing a common database for data from all of your UNIX systems
- o Retaining UNIX system profile information
- o Producing reports that enable you to compare data across UNIX systems and over time

The Analyzer helps you make the best use of computing resources by:

- o Reporting on CPU usage and bottlenecks, and providing a detailed system performance summary
- o Retaining data that you can use to develop your own reports

The Analyzer helps you analyze availability, response time, and throughput by:

- o Reporting on LAN and disk traffic and transaction response
- o Retaining data that you can use to develop your own reports

Analyzing UNIX Measurements

The tables in this chapter help you identify CA MICS elements useful for analyzing data by business area.

To analyze these measures ...	Use this table
Common Metrics	2-0
Global System	2-1
Application	2-2
Killed Process	2-3
Process Exception	2-4
LAN Activity	2-5
Disk Activity	2-6
Logical Volume	2-7
Transaction Tracker	2-8
System Configuration & Collector Profile	2-9

COMMON METRICS

HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
RECORD_TYPE	RTYPE	All Releases	All Releases	All Releases
DATE_SECOND	WRKTS	All Releases	All Releases	All Releases
INTERVAL	DURATION	All Releases	All Releases	All Releases
APP_SAMPLE	AUMSAMPL	All Releases	All Releases	All Releases
GBL_COLLECTOR	AUMSCOPE	All Releases	All Releases	All Releases
PROC_TTY	AUMLNTTY	All Releases	All Releases	All Releases
GBL_OSNAME	AUMOPSYS	All Releases	All Releases	All Releases
GBL_OSRELEASE	AUMOPSVR	All Releases	All Releases	All Releases
GBL_OSVERSION	AUMOPSLC	All Releases	All Releases	All Releases
GBL_MACHINE	AUMCPMDL	All Releases	All Releases	All Releases
APP_NAME	AUMAPU	All Releases	All Releases	All Releases
GBL_LOGFILE_VERSION	AUMLGVER	All Releases	All Releases	All Releases
APP_NUM	AUMAPNUM	All Releases	All Releases	All Releases
BYDSK_ID	AUMLDEV	All Releases	n/a	n/a
PROC_PROC_ID	AUMPIN	All Releases	All Releases	All Releases
PROC_PARENT_PROC_ID	AUMPPID	All Releases	All Releases	All Releases
PROC_GROUP_ID	AUMGRPID	All Releases	All Releases	All Releases
PROC_PRI	AUMPRI	All Releases	All Releases	n/a
PROC_PROC_NAME	AUMPROG	All Releases	All Releases	All Releases
PROC_USER_NAME	AUMUSER	All Releases	All Releases	All Releases
PROC_APP_ID	AUMAPNUM	All Releases	All Releases	All Releases
FS_DIRNAME	AUMDIRNM	C.02.60 and up	n/a	n/a

Table 2-0. Common Metrics

GLOBAL METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
CPU Use					
	GBL_CPU_TOTAL_TIME	GPRCPUTM	All Releases	All Releases	All Releases
	GBL_CPU_USER_MODE_TIME	GPRUSRTM	All Releases	All Releases	All Releases
	GBL_CPU_REALTIME_TIME	GPRRTMTM	All Releases	n/a	n/a
	GBL_CPU_NICE_TIME	GPRNICTM	All Releases	n/a	n/a
	GBL_CPU_NORMAL_TIME	GPRNORTM	All Releases	n/a	n/a
	GBL_CPU_SYS_MODE_TIME	GPRSYSTM	All Releases	All Releases	All Releases
	GBL_CPU_CSWITCH_TIME	GPRCSWTM	All Releases	n/a	n/a
	GBL_CPU_INTERRUPT_TIME	GPRINTTM	All Releases	n/a	n/a
	GBL_CPU_SYSCALL_TIME	GPRSSCTM	All Releases	All Releases	All Releases
	GBL_CPU_WAIT_TIME	GPRPAUTM	n/a	All Releases	All Releases
	GBL_ACTIVE_CPU	GPRACCPU	C.02.65 and up	C.04.5 and up	C.02.05 and up
	** CALCULATED **	GPRCPUCP	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCCPU	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCUSR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCNIC	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCNOR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCRTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCSYS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCCSW	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCINT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCSSC	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRNPROC	Check Derivation	Check Derivation	Check Derivation
Memory Use					
	GBL_MEM_SWAP	GPRSWPCT	All Releases	n/a	n/a
	GBL_MEM_PAGE_REQUEST	GPRPFLCT	All Releases	All Releases	All Releases
	GBL_MEM_PAGEOUT	GPRPGOCT	All Releases	All Releases	All Releases
	** CALCULATED **	GPRMMEMS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMMEMU	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMEMUT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSWAPU	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRACTVM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCHIT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSWOCT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCMEM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCMMS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCMMU	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCAVM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPCSWS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSWPRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSWORT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPAGRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSHMSG	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSMEMS	Check Derivation	Check Derivation	Check Derivation

Table 2-1. Global System Data (Part 1 of 4)

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Disk Use					
	GBL_DISK_PHYS_READ	GPRPRDCT	All Releases	All Releases	All Releases
	GBL_DISK_PHYS_WRITE	GPRPWTCT	All Releases	All Releases	All Releases
	GBL_DISK_PHYS_IO	GPRDSKIO	All Releases	All Releases	All Releases
	GBL_DISK_FS_READ	GPRDFRCT	All Releases	n/a	n/a
	GBL_DISK_FS_WRITE	GPRDFWCT	All Releases	n/a	n/a
	GBL_DISK_FS_IO	GPRFSDCT	All Releases	n/a	n/a
	GBL_DISK_VM_READ	GPRDMRCT	All Releases	All Releases	n/a
	GBL_DISK_VM_WRITE	GPRDMWCT	All Releases	All Releases	n/a
	GBL_DISK_VM_IO	GPRDMMCT	All Releases	All Releases	All Releases
	GBL_DISK_SYSTEM_READ	GPRDSRCT	All Releases	n/a	n/a
	GBL_DISK_SYSTEM_WRITE	GPRDSWCT	All Releases	n/a	n/a
	GBL_DISK_SYSTEM_IO	GPRFSMCT	All Releases	n/a	n/a
	GBL_DISK_RAW_READ	GPRDRRCT	All Releases	All Releases	All Releases
	GBL_DISK_RAW_WRITE	GPRDRWCT	All Releases	All Releases	All Releases
	GBL_DISK_RAW_IO	GPRRAWCT	All Releases	All Releases	All Releases
	GBL_DISK_BLOCK_READ	GPRDBRCT	n/a	All Releases	All Releases
	GBL_DISK_BLOCK_WRITE	GPRDBWCT	n/a	All Releases	All Releases
	GBL_DISK_BLOCK_IO	GPRDBKCT	n/a	All Releases	All Releases
	GBL_DISK_LOGL_READ	GPRLRDCT	All Releases	n/a	n/a
	GBL_DISK_LOGL_WRITE	GPRLWTCT	All Releases	n/a	n/a
	GBL_DISK_LOGL_IO	GPRLDKCT	All Releases	n/a	n/a
	GBL_DISK_PHYS_BYTE	GPRPKBCT	All Releases	All Releases	All Releases
	GBL_DISK_LOGL_READ_BYTE	GPRLRDKB	All Releases	n/a	n/a
	GBL_DISK_LOGL_WRITE_BYTE	GPRLWTKB	All Releases	n/a	n/a
	GBL_DISK_UTIL_PEAK	GPRMPBDU	All Releases	All Releases	All Releases
	GBL_DISK_TIME_PEAK	GPRMXDTM	All Releases	All Releases	All Releases
	GBL_FS_SPACE_UTIL_PEAK	GPRMPDSU	All Releases	All Releases	All Releases
	** CALCULATED **	GPRRDHIT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRDSKRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRDMRRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRDMWRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRDMMRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRDKLRT	Check Derivation	Check Derivation	Check Derivation
Networking					
	GBL_NET_IN_PACKET	GPRLPCKI	All Releases	C.02.58 and up	C.02.05 and up
	GBL_NET_OUT_PACKET	GPRLPCKO	All Releases	C.02.58 and up	C.02.05 and up
	GBL_NFS_CALL	GPRNFSCT	All Releases	All Releases	All Releases
	** CALCULATED **	GPRLCOLS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRLERRS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVLCL	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVLER	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRNFSRT	Check Derivation	Check Derivation	Check Derivation

Table 2-1. Global System Data (Part 2 of 4)

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
User/Process Count					
	GBL_COMPLETED_PROC	GPRCPRCT	All Releases	All Releases	All Releases
	GBL_STARTED_PROC	GPRSPRCT	All Releases	All Releases	All Releases
	GBL_QUEUE_HISTOGRAM	GPRWAITQ	C.00.01 and up	All Releases	ALl Releases
	** CALCULATED **	GPRAVAPR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVPRO	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVUSR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVBIQ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVRTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRRUNTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPROCT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAPRCT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRUSRCT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPTBLU	Check Derivation	Check Derivation	Check Derivation
Queue Length					
	GBL_NET_OUTQUEUE	GPROUTQU	C.02.00 and up	n/a	n/a
	** CALCULATED **	GPRAVCPQ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVPRQ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVMMQ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVDKQ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVIPQ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVNTQ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVSLQ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVSUS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSUSCT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRAVSWA	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSWACT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRCPQTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRPRITM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMMQTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRDKQTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRIPQTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRNTQTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSLQTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMQUSE	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMSGQU	Check Derivation	Check Derivation	Check Derivation
Identifying					
	GBL_PROC_SAMPLE	AUMSAMPL	All Releases	All Releases	All Releases

Table 2-1. Global System Data (Part 3 of 4)

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
System Configuration			Check Derivation	Check Derivation	Check Derivation
			Check Derivation	Check Derivation	Check Derivation
	TBL_BUFFER_CACHE_AVAIL	GPRCACHE	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMXCPU	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMXDSK	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMXLAN	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRCACHE	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMEMSW	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMEMSZ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRMEMUS	Check Derivation	Check Derivation	Check Derivation
System and Collector					
	GBL_SYSTEM_UPTIME_HOURS	GPRSUPTM	C.02.00 and up	n/a	C.02.05 and up
	GBL_TT_OVERFLOW_COUNT	GPROFLOW	C.02.00 and up	n/a	C.02.53 and up
	GBL_LOST_MI_TRACE_BUFFERS	GPRLOSTB	C.02.00 and up	n/a	C.02.53 and up
	TBL_INODE_CACHE_USED	GPRINODU	C.02.00 and up	n/a	C.02.53 and up
	TBL_BUFFER_CACHE_USED	GPRCHUSE	C.02.00 and up	n/a	n/a
	TBL_FILE_LOCK_USED	GPRLUSED	n/a	n/a	C.02.53 and up
Other Measurements					
	** CALCULATED **	GPRSCALL	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSCLRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSEMID	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRSEMST	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRTBLEN	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	GPRTBLUK	Check Derivation	Check Derivation	Check Derivation

Table 2-1. Global System Data (Part 4 of 4)

APPLICATION METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
CPU Use					
	APP_CPU_TOTAL_TIME	APPCPUTM	All Releases	All Releases	All Releases
	APP_CPU_USER_MODE_TIME	APPUSRTM	n/a	All Releases	All Releases
	APP_CPU_NICE_TIME	APPNICTM	All Releases	n/a	n/a
	APP_CPU_NORMAL_TIME	APPNORTM	All Releases	n/a	n/a
	APP_CPU_REALTIME_TIME	APPRTMTM	All Releases	n/a	n/a
	APP_CPU_SYS_MODE_TIME	APPSYSTM	All Releases	All Releases	All Releases
	** CALCULATED **	APPPCCPU	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPCUSR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPCNIC	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPCNOR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPCRTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPCSYS	Check Derivation	Check Derivation	Check Derivation
Memory Use					
	APP_MEM_VIRT	APPMEMVR	All Releases	All Releases	All Releases
	APP_MEM_RES	APPMEMRS	All Releases	n/a	n/a
	APP_MEM_UTIL	APPMEMOR	n/a	All Releases	All Releases
	** CALCULATED **	APPMEMPR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPCMPR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPCMRS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPSWAPU	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPFMAJ	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPFMIN	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPMAJRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPMINRT	Check Derivation	Check Derivation	Check Derivation
Wait States					
	APP_PRI_WAIT_PCT	APPPCSCP	All Releases	n/a	n/a
	APP_DISK_SUBSYSTEM_WAIT_PCT	APPPCSDK	All Releases	n/a	n/a
	APP_MEM_WAIT_PCT	APPPCSWP	All Releases	n/a	n/a
	APP_SEM_WAIT_PCT	APPPCSIM	All Releases	n/a	n/a
	APP_IPC_SUBSYSTEM_WAIT_PCT	APPPCSIP	All Releases	n/a	n/a
	APP_NETWORK_SUBSYSTEM_ WAIT_PCT	APPPCSLN	All Releases	n/a	n/a
	APP_OTHER_IO_WAIT_PCT	APPPCSOT	All Releases	n/a	n/a
	APP_SLEEP_WAIT_PCT	APPPCSDR	All Releases	n/a	n/a

Table 2-2. Application Data (Part 1 of 3)

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Process Cnt					
	APP_COMPLETED_PROC	APPCPRCT	All Releases	All Releases	All Releases
	** CALCULATED **	APPAVAPR	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPAVPRO	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPAVSUS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPRUNTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPAVRTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPAPRCT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPROCT	Check Derivation	Check Derivation	Check Derivation
Disk Use					
	APP_DISK_FS_IO	APPDFSCT	All Releases	n/a	n/a
	APP_DISK_VM_IO	APPDMMCT	All Releases	n/a	n/a
	APP_DISK_RAW_IO	APPDRAIO	All Releases	n/a	n/a
	APP_DISK_SYSTEM_IO	APPDSYIO	All Releases	n/a	n/a
	APP_DISK_PHYS_IO	APPDSKIO	All Releases	All Releases	C.02.05 and up
	APP_DISK_LOGL_READ	APPLRDCT	All Releases	n/a	n/a
	APP_DISK_LOGL_WRITE	APPLWTCT	All Releases	n/a	n/a
	APP_DISK_LOGL_IO	APPLIOCT	All Releases	n/a	n/a
	APP_IO_BYTE	APPPKBCT	All Releases	All Releases	C.02.05 and up
	APP_DISK_PHYS_READ	APPPRDCT	All Releases	n/a	n/a
	APP_DISK_PHYS_WRITE	APPPWTCT	All Releases	n/a	n/a
	APP_DISK_BLOCK_IO	APPDBKCT	n/a	All Releases	All Releases
	APP_DISK_BLOCK_READ	APPDBRCT	n/a	All Releases	All Releases
	APP_DISK_BLOCK_WRITE	APPDBWCT	n/a	All Releases	All Releases
	** CALCULATED **	APPDFSRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPDMMRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPDRART	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPDSYRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPKBRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPSPRD	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	APPPSPWT	Check Derivation	Check Derivation	Check Derivation
Process Res. Manager					
	APP_PRM_LOGGING_MODE	APPLOGMD	C.03.20 and up	n/a	n/a
	APP_PRM_MEM_ENTITLEMENT	APPMEMEN	C.03.20 and up	n/a	n/a
	APP_PRM_MEM_UPPERBOUND	APPMEMUP	C.03.20 and up	n/a	n/a
	APP_PRM_MEM_STATE	APPMEMST	C.03.20 and up	n/a	n/a
	APP_PRM_STATE	APPCPUST	C.03.20 and up	n/a	n/a
	APP_PRM_MEM_AVAIL	APPMEMAV	C.03.20 and up	n/a	n/a
	APP_PRM_CPU_ENTITLEMENT	APPCPUEN	C.03.20 and up	n/a	n/a

Table 2-2. Application Data (Part 2 of 3)

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Other Meas.					
	** EXIT **	APPCOST			
PRM Use					
	APP_PRM_CPUCAP_MODE	APPCPUCM	All Releases	n/a	n/a
	APP_PRM_CPU_ENTITLEMENT	APPCPUEN	C.03.20 and up	n/a	n/a
	APP_PRM_STATE	APPCPUST	C.03.20 and up	n/a	n/a
	APP_PRM_LOGGING_MODE	APPLOGMD	C.03.20 and up	n/a	n/a
	APP_PRM_MEM_AVAIL	APPMEMAV	C.03.20 and up	n/a	n/a
	APP_PRM_MEM_ENTITLEMENT	APPMEMEN	C.03.20 and up	n/a	n/a
	APP_PRM_MEM_STATE	APPMEMST	C.03.20 and up	n/a	n/a
	APP_PRM_MEM_UPPERBOUND	APPMEMUP	C.03.20 and up	n/a	n/a
	** CALCULATED **	APPMEMUS	Check Derivation	Check Derivation	Check Derivation

Table 2-2. Application Data (Part 3 of 3)

KILLED PROCESS METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
CPU Use					
	PROC_CPU_TOTAL_TIME_CUM	KPSCPUTM	All Releases	All Releases	All Releases
	** CALCULATED **	KPSRUNTM	Check Derivation	Check Derivation	Check Derivation
Memory Use					
	** CALCULATED **	KPSRSSKB	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	KPSVSSKB	Check Derivation	Check Derivation	Check Derivation
Disk Use					
	PROC_DISK_PHYS_IO_CUM	KPSDSKIO	All Releases	n/a	n/a
	PROC_DISK_BLOCK_IO_CUM	KPSDBKCT	n/a	All Releases	C.03.45 and up
	PROC_IO_BYTE_CUM	KPSPKBCT	All Releases	All Releases	C.03.45 and up
	PROC_DISK_LOGL_IO_CUM	KPSLDKCT	All Releases	n/a	n/a
	** CALCULATED **	KPSAVDRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	KPSDBKRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	KPSPKBRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	KPSAVLRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	KPSSDKTM	Check Derivation	Check Derivation	Check Derivation

Table 2-3. Killed Process Data (Part 1 of 2)

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Wait States					
	PROC_PRI_WAIT_TIME	KPSSCPTM	All Releases	n/a	n/a
	PROC_SYS_WAIT_TIME	KPSSSYTM	All Releases	n/a	n/a
	PROC_MEM_WAIT_TIME	KPSSVMTM	All Releases	n/a	n/a
	PROC_SLEEP_WAIT_TIME	KPSSDRTM	All Releases	n/a	n/a
	PROC_SEM_WAIT_TIME	KPSSIMTM	All Releases	n/a	n/a
	PROC_IPC_SUBSYSTEM_WAIT_ TIME	KPSSIPTM	All Releases	n/a	n/a
	PROC_LAN_WAIT_TIME	KPSSLNTM	All Releases	n/a	n/a
	PROC_NFS_WAIT_TIME	KPSSNFTM	All Releases	n/a	n/a
	PROC_OTHER_IO_WAIT_TIME	KPSSOITM	All Releases	n/a	n/a

Table 2-3. Killed Process Data (Part 2 of 2)

PROCESS EXCEPTION METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
CPU Use					
	PROC_CPU_TOTAL_TIME	IPSCPUTM	All Releases	All Releases	All Releases
	PROC_CPU_USER_MODE_TIME	IPSUSRTM	All Releases	All Releases	All Releases
	PROC_CPU_NICE_TIME	IPSNICTM	All Releases	n/a	n/a
	PROC_CPU_NORMAL_TIME	IPSNORTM	All Releases	n/a	n/a
	PROC_CPU_REALTIME_TIME	IPSRTMTM	All Releases	n/a	n/a
	PROC_CPU_SYS_MODE_TIME	IPSSYSTM	All Releases	All Releases	All Releases
	PROC_CPU_CSWITCH_TIME	IPSCSWTM	All Releases	n/a	n/a
	PROC_CPU_INTERRUPT_TIME	IPSINTTM	All Releases	n/a	n/a
	PROC_CPU_SYSCALL_TIME	IPSSSCTM	All Releases	n/a	n/a
	PROC_RUN_TIME	IPSRUNTM	All Releases	All Releases	All Releases
	PROC_INTERVAL_ALIVE	IPSALITM	All Releases	All Releases	All Releases
	PROC_CPU_TOTAL_TIME_CUM	IPSCUMTM	Check Derivation	Check Derivation	Check Derivation
Memory Use					
	PROC_MAJOR_FAULT	IPSPFMAJ	All Releases	All Releases	All Releases
	PROC_MINOR_FAULT	IPSPFMIN	All Releases	All Releases	All Releases
	PROC_MEM_RES	IPSRSSKB	All Releases	All Releases	All Releases
	PROC_MEM_VIRT	IPSVSSKB	All Releases	All Releases	All Releases
Disk Use					
	PROC_DISK_PHYS_IO	IPSDSKIO	All Releases	n/a	n/a
	PROC_DISK_FS_READ	IPSDFRCT	All Releases	n/a	n/a
	PROC_DISK_FS_WRITE	IPSDFWCT	All Releases	n/a	n/a
	PROC_DISK_FS_IO	IPSFSDCT	All Releases	n/a	n/a
	PROC_DISK_VM_IO	IPSDMMCT	All Releases	n/a	n/a
	PROC_DISK_SYSTEM_IO	IPSDSYCT	All Releases	n/a	n/a
	PROC_IO_BYTE	IPSPKBCT	All Releases	All Releases	C.03.45 and up
	PROC_DISK_BLOCK_IO	IPSDBKCT	n/a	All Releases	All Releases
	PROC_DISK_BLOCK_READ	IPSDBRCT	n/a	All Releases	All Releases
	PROC_DISK_BLOCK_WRITE	IPSDBWCT	n/a	All Releases	All Releases
	PROC_DISK_LOGL_READ	IPSLRDCT	All Releases	n/a	n/a
	PROC_DISK_LOGL_WRITE	IPSLWTCT	All Releases	n/a	n/a
Identifying Metrics					
	PROC_INTEREST	IPSINTR	All Releases	All Releases	All Releases
	PROC_STOP_REASON	IPSSTOPR	All Releases	All Releases	All Releases
	PROC_PRMID	IPSPRMID	C.03.25 and up	n/a	n/a
	PROC_THREAD_COUNT	IPSTHCNT	C.02.00 and up	n/a	n/a

Table 2-4. Process Exception Data

NETWORK INTERFACE (LAN) METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Errors and Collisions					
	BYNETIF_COLLISION	NIFCOLCT	All Releases	All Releases	All Releases
	BYNETIF_ERROR	NIFERRCT	All Releases	All Releases	All Releases
	** CALCULATED **	NIFPSCOL	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	NIFPSERR	Check Derivation	Check Derivation	Check Derivation
Traffic					
	BYNETIF_IN_PACKET	NIFPKICT	All Releases	All Releases	All Releases
	BYNETIF_OUT_PACKET	NIFPKOCT	All Releases	All Releases	All Releases
	BYNETIF_NET_SPEED	NIFSPEED	C.03.25 and up	C.03.40	n/a
	BYNETIF_NET_MTU	NIFMTUSZ	C.03.25 and up	n/a	n/a
	** CALCULATED **	NIFPSPKI	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	NIFPSPKO	Check Derivation	Check Derivation	Check Derivation
Identifying					
	BYNETIF_NAME	NIFNAME	All Releases	All Releases	All Releases
	BYNETIF_QUEUE	NIFQUEUE	C.03.70 and up	n/a	n/a

Table 2-5. LAN Activity Data

DISK METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Disk Busy					
	** CALCULATED **	DDAPCDUT	Check Derivation	Check Derivation	Check Derivation
I/O					
	BYDSK_PHYS_IO	DDADSKIO	All Releases	All Releases	All Releases
	BYDSK_PHYS_READ	DDAPRDCT	All Releases	All Releases	All Releases
	BYDSK_PHYS_WRITE	DDAPWTCT	All Releases	All Releases	All Releases
	BYDSK_PHYS_BYTE	DDAPKBCT	All Releases	All Releases	All Releases
	BYDSK_PHYS_READ_BYTE	DDARKBCT	n/a	All Releases	All Releases
	BYDSK_PHYS_WRITE_BYTE	DDAWKBCT	n/a	All Releases	C.02.05 and up
	BYDSK_FS_READ	DDADFRCT	All Releases	n/a	n/a
	BYDSK_FS_WRITE	DDADFWCT	All Releases	n/a	n/a
	BYDSK_VM_IO	DDADMMCT	All Releases	n/a	n/a
	BYDSK_SYSTEM_IO	DDASYSCT	All Releases	n/a	n/a
	BYDSK_RAW_READ	DDADRRCT	All Releases	n/a	n/a
	BYDSK_RAW_WRITE	DDADRWCT	All Releases	n/a	n/a
	BYDSK_LOGL_READ	DDALRDCT	All Releases	n/a	n/a
	BYDSK_LOGL_WRITE	DDALWTCT	All Releases	n/a	n/a
	** CALCULATED **	DDADSKRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPKBRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAWKBRT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPSFRD	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPSFWT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPSMEM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPSSYS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPSRRD	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPSRWT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPSLRD	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAPSLWT	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDADDPDR	Check Derivation	Check Derivation	Check Derivation
Queue Length					
	** CALCULATED **	DDAAVQLN	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	DDAQLNSM	Check Derivation	Check Derivation	Check Derivation
Identifying					
	BYDSK_DEVNAME	DDANAME	All Releases	All Releases	All Releases
	BYDSK_DIRNAME	DDADIR	All Releases	n/a	n/a

Table 2-6. Disk Activity Data

LOGICAL VOLUME METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Space Use					
	** CALCULATED **	VLMPCSPU	Check Derivation	n/a	n/a
	** CALCULATED **	VLMSPUSM	Check Derivation	n/a	n/a
	** CALCULATED **	VLMMNSPU	Check Derivation	n/a	n/a
	** CALCULATED **	VLMMXSPU	Check Derivation	n/a	n/a
I/O					
	LV_LOGL_READ	VLMLRDCT	up to C.01.00	n/a	n/a
	LV_LOGL_WRITE	VLMLWTCT	up to C.01.00	n/a	n/a
	** CALCULATED **	VLMPSLRD	Check Derivation	n/a	n/a
	** CALCULATED **	VLMPSLWT	Check Derivation	n/a	n/a
Identifying					
	LV_GROUP_NAME	VLMNAME	All Releases	n/a	C.02.06 and up
	LV_DIRNAME	VLMDIR	All Releases	n/a	C.02.06 and up
	FS_DIRNAME	VLMPTHNM	All Releases	n/a	n/a

Table 2-7. Logical Volume Data

TRANSACTION METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Parameter Settings					
	TT_SLO_THRESHOLD	TTASLOMX	All Releases	All Releases	All Releases
	TT_NUM_BINS	TTABINCT	All Releases	All Releases	All Releases
	TTBIN_UPPER_RANGE_(1-10)	TTAUPRxx	All Releases	All Releases	All Releases
Completed Transaction Activity					
	TT_COUNT	TTATRANS	All Releases	All Releases	All Releases
	TTBIN_TRANS_COUNT_(1-10)	TTATRNxx	All Releases	All Releases	All Releases
	TT_SLO_COUNT	TTAVSLCT	All Releases	All Releases	All Releases
	** CALCULATED **	TTATRNTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	TTAPCBXx	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	TTAAVTTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	TTAPCVSL	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	TTALGLIO	Check Derivation	Check Derivation	Check Derivation
Aborted Tran Activity					
	TT_ABORT	TTAABORT	All Releases	All Releases	All Releases
	TT_FAILED	TTATFAIL	C.01.00 and up	C.02.58 and up	C.02.05 and up
	** CALCULATED **	TTAABTTM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	TTAAVATM	Check Derivation	Check Derivation	Check Derivation
Correlator Metrics					
	TT_CLIENT_ADDRESS	TTACORAD	C.02.00 and up	C.02.58 and up	C.02.05 and up
	TT_CLIENT_ADDRESS_FORMAT	TTAADFMT	C.02.00 and up	C.02.58 and up	C.02.05 and up
	TT_CLIENT_TRAN_ID	TTATRNID	C.02.00 and up	C.02.58 and up	C.02.05 and up
ARM Use					
	TT_TRAN_ID	TTAARMID	All Releases	C.02.58 and up	All Releases
	TT_INFO	TTAARMIF	C.01.00 and up	C.02.58 and up	C.02.05 and up
	TT_APP_NAME	TTAARMNM	C.01.00 and up	C.02.58 and up	C.02.05 and up
	TT_UNAME	TTAARMUN	C.01.00 and up	C.02.58 and up	C.02.05 and up
	TT_APP_TRAN_NAME	TTAATNM	C.01.00 and up	C.02.58 and up	All Releases
CPU Use					
	** CALCULATED **	TTACPUTM	Check Derivation	Check Derivation	Check Derivation
Other					
	TT_NAME	TTANAME	All Releases	All Releases	All Releases
	** CALCULATED **	TTADSKIO	Check Derivation	Check Derivation	Check Derivation

Table 2-8. Business Transaction SLO Activity Data

CONFIGURATION METRICS

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
CPU Use					
	GBL_NUM_CPU	MCFNCPUS	All Releases	All Releases	All Releases
	TBL_PROC_TABLE_AVAIL	MCFNPROC	All Releases	All Releases	All Releases
	** CALCULATED **	MCFMNCPU	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMXCPU	Check Derivation	Check Derivation	Check Derivation
Memory Configuration					
	TBL_BUFFER_CACHE_AVAIL	MCFCACHE	All Releases	All Releases	All Releases
	GBL_SWAP_SPACE_AVAIL_KB	MCFMEMSW	All Releases	All Releases	All Releases
	GBL_MEM_PHYS	MCFMEMSZ	All Releases	All Releases	All Releases
	GBL_MEM_AVAIL	MCFMEMUS	All Releases	All Releases	All Releases
	TBL_SHMEM_TABLE_AVAIL	MCFSHMSG	C.02.00 and up	n/a	C.02.53 and up
	GBL_THRESHOLD_PROCMEM	MCFPRCMM	C.02.40 and up	n/a	C.03.70 and up
	** CALCULATED **	MCFMNCHC	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMNMEM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMNMUS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMNSWS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMXCHC	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMXMEM	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMXMUS	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMXSWS	Check Derivation	Check Derivation	Check Derivation
Disk Configuration					
	GBL_NUM_DISK	MCFNDSKS	All Releases	All Releases	All Releases
	** CALCULATED **	MCFMNDSK	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMXDSK	Check Derivation	Check Derivation	Check Derivation
LAN Configuration					
	GBL_NUM_NETWORK	MCFNLANS	All Releases	All Releases	All Releases
	** CALCULATED **	MCFMNLAN	Check Derivation	Check Derivation	Check Derivation
	** CALCULATED **	MCFMXLAN	Check Derivation	Check Derivation	Check Derivation

Table 2-9. System Configuration & Collector Profile Data (Part 1 of 2)

Metric Type	HP Metric	CA MICS Data Element	HP-UX Platform OpenView Release	AIX Platform OpenView Release	SOL Platform OpenView Release
Collector Setting					
	GBL_THRESHOLD_CPU	MCFCPUMN	All Releases	All Releases	All Releases
	GBL_THRESHOLD_DISK	MCFIORMN	All Releases	All Releases	All Releases
	GBL_THRESHOLD_SHORTLIVED	MCFLOSHR	All Releases	n/a	n/a
	GBL_THRESHOLD_NOKILLED	MCFNOKIL	All Releases	All Releases	All Releases
	GBL_THRESHOLD_NONEW	MCFNONEW	All Releases	All Releases	All Releases
	GBL_THRESHOLD_WAIT_CPU	MCFWCPMN	All Releases	n/a	n/a
	GBL_THRESHOLD_WAIT_DISK	MCFWDKMN	All Releases	n/a	n/a
	GBL_THRESHOLD_WAIT_MEMORY	MCFWMEMN	All Releases	n/a	n/a
	GBL_THRESHOLD_WAIT_IMPEDE	MCFWIMMN	All Releases	n/a	n/a
	GBL_LOGGING_TYPES	MCFLGTPS	All Releases	All Releases	All Releases
	GBL_SYSTEM_ID	AUMSYSID	All Releases	All Releases	All Releases
	GBL_MACHINE_MODEL	MCFMCMOD	C.02.40 and up	n/a	n/a
	GBL_OSKERNELTYPE_INT	MCFKRNSZ	C.02.00 and up	n/a	n/a
	TBL_FILE_TABLE_AVAIL	MCFTBLUK	C.02.00 and up	n/a	C.02.53 and up
	TBL_FILE_LOCK_AVAIL	MCFLCKAV	C.02.00 and up	n/a	n/a
	TBL_INODE_CACHE_AVAIL	MCFINODE	C.02.00 and up	n/a	C.02.53 and up
	TBL_SEM_TABLE_AVAIL	MCFSEMID	C.02.00 and up	n/a	C.02.53 and up
	TBL_MSG_TABLE_AVAIL	MCFMSGQU	C.02.00 and up	n/a	C.02.53 and up
	GBL_THRESHOLD_PROCMEM	MCFPRCMM	C.02.40 and up	n/a	C.03.75 and up
	GBL_SUBPROCSAMPLEINTERVAL	MCFPRCIN	n/a	C.02.58 and up	C.02.53 and up

Table 2-9. System Configuration & Collector Profile Data (Part 2 of 2)

2.2 - CA MICS Product Interfaces

CA MICS Accounting and Chargeback

With ACT6260, CA MICS Accounting and Chargeback will support the following MW OS Analyzer data elements that you can use for billing resource or transaction usage. The elements that can be billed are:

- UNIX/MW OS Application Total CPU Time
- UNIX/MW OS Application System CPU Time
- UNIX/MW OS Application User CPU Time
- UNIX/MW OS Application Nice CPU Time
- UNIX/MW OS Application Real Time CPU Time
- UNIX/MW OS Application CPU Use at Normal Priorities
- UNIX/MW OS Application Total Physical Disk I/O
- UNIX/MW OS Application Logical Read Disk I/O
- UNIX/MW OS Application Logical Write Disk I/O
- UNIX/MW OS Application Physical Read Disk I/O
- UNIX/MW OS Application Physical Write Disk I/O
- UNIX/MW OS Application I/O Transfer Cnt in Kb
- UNIX/MW OS Process CPU Time
- UNIX/MW OS Process Total Physical Disk I/O
- UNIX/MW OS Process Logical Disk I/O
- UNIX/MW OS Process I/O Transfer Cnt in Kb

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 3

REPORTS

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 3

REPORTS

CHAPTER 3. MW OS AGENT ANALYZER REPORTS

The MW OS Agent Analyzer produces reports using the CA MICS Information Center Facility (MICF), which is part of the CA MICS Workstation Facility (MWF).

MICF inquiries produce meaningful reports from the CA MICS database and provide you with the flexibility to code and save your own report formats. You can execute and easily modify distributed MICF inquiries using the Database Inquiries option from the MICF main menu.

Refer to the CA MICS MICF Reference Guide for more information on MICF inquiries.

The following topic is presented:

1 - MICF Inquiries for the MW OS Agent Analyzer

3.1 - MICF Inquiries for the MW OS Agent Analyzer

The table below lists the MICF inquiries distributed with the MW OS Agent Analyzer. The remainder of this chapter follows this category listing. Each section contains a brief description of the inquiry and the panels used to generate it.

The Analyzer's color graphic reports (inquiry ID AUMCxx) cannot be reproduced in this guide, however the line graphics (inquiry ID AUMLxx) are included.

Category	Report Name	Inquiry ID	Use	For additional information, use ...
Multiple Systems	UNIX Global CPU Utilization	AUMCS1	Overview of systems and their use	n/a
	UNIX Global Queue Depth Comparison	AUMCS2	Overview of systems and their bottlenecks	n/a
Single Systems	UNIX Global Bottlenecks	AUMCS4	Identify bottlenecks	n/a
	UNIX Global Disk Summary	AUMCS6	Identify number of disk I/Os and busiest disk.	AUMCSB and AUMCSC
Applica-tions	UNIX Application Performance Summary	AUMLS2	Detailed application information.	AUMLS3
Killed Processes	UNIX Detail Killed Process	AUMLS3	Identify programs that make up the application group.	n/a
Disk	UNIX Top 5 Disk Utilization Comparison	AUMCSB	Identify busiest disks.	AUMCSC
	Disk I/O Activity	AUMCSC	Identify how a disk is being used.	n/a
	UNIX Logical Volume Activity	AUMCSD	Identify logical volume logical I/Os and disk space usage.	n/a
Global	UNIX System Performance Summary	AUMLS1	Identify how the system is being used.	n/a

The following sections describe the individual inquiries:

1 - UNIX Multiple System Inquiries
2 - UNIX Single System Inquiries
3 - UNIX Application Inquiries
4 - UNIX Process Activity Inquiries
5 - UNIX Disk Inquiries
6 - UNIX Global Inquiries

3.1.1 - UNIX Multiple System Inquiries

The MW OS Agent Analyzer provides the following inquiries that report on data from multiple UNIX systems:

1 - UNIX Global CPU Utilization (AUMCS1)
2 - UNIX Global Queue Depth Compare (AUMCS2)

3.1.1.1 - UNIX Global CPU Utilization (AUMCS1)

The UNIX Global CPU Utilization chart uses measurements from the CA MICS AUMGPR file to display an overall view of processor utilization for multiple systems over the date and/or time ranges you select. The color graphic box charts display and categorize each system's utilization by zone and/or date intervals based on the report summarization options you select.

Each box on the chart depicts processor utilization over time or date. Multiple systems are displayed on separate rows of the chart. Actual processor busy is displayed below each box. The type of utilization activity is denoted by color within the boxes, where utilization activity types are:

For HP-UX systems:
- o Context switch (GPRPCCSW)
- o Normal (GPRPCNOR)
- o Interrupt (GPRPCINT)
- o Real time (GPRPCRTM)
- o Nice (GPRPNOR)
- o System calls (GPRPCSSC)

For Solaris or AIX systems:
- o User mode (GPRPCUSR)
- o System mode (GPRPCSYS)

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMCS1, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result displays on your terminal.

```
---------------------- UNIX Global CPU Utilization ----------------------
Command ===>

Composing CA MICS Inquiry:

Database ===> ____________________________________________
Cycle(s) ===> __ - __
Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
Plot by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
Generate a plot for each value of Zone     ===> _ (Y/N)

Data Selection:
                         Date Range 1         Date Range 2
                     dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
  Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
  Hour         ===> __ - __    __ - __    __ - __
  Zone         ===> _ _ _ _ _ _ _ _ _ _
  Sysid        ===> ____ ____ ____ ____ ____ ____
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Plot by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY level. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a plot for each value of Zone - Required. Defaults to Y. Reports for each zone can be generated when requested if you specify Y (Yes). The nature and resource requirements of UNIX processing very often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as 01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

3.1.1.2 - UNIX Global Queue Depth Comparison (AUMCS2)

The UNIX Global Queue Depth Comparison chart uses measurements from the CA MICS AUMGPR file to display average process wait queue metrics for multiple systems over the date and/or time ranges you select. This color graphic box chart displays and categorizes each system's queue depth by zone and/or date intervals based on the report summarization options you select.

Each box on the chart depicts system queue depth over time or date. Multiple systems are displayed on separate rows of the chart. Total average queue depth is shown below each box. The type of resource causing the wait queue is denoted by color within the boxes, where the resource types are disk (GPRAVDKQ), network (GPRAVNTQ), IP communications (GPRAVIPQ), memory (GPRAVMMQ), run (GPRAVCPQ), and priority (GPRAVPRQ).

This report is not supported for Solaris or AIX systems because the elements needed are not all available from the data source.

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMCS2, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result displays on your terminal.

```
---------- UNIX Global Queue Depth Comparison ----------
Command ===>

Composing CA MICS Inquiry:

Database ===> ____________________________________________
Cycle(s) ===> __ - __
Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
Plot by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
Generate a plot for each value of Zone     ===> _ (Y/N)

Data Selection:
                        Date Range 1         Date Range 2
                     dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
  Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
  Hour         ===> __ - __    __ - __    __ - __
  Zone         ===> _ _ _ _ _ _ _ _ _ _
  Sysid        ===> ____ ____ ____ ____ ____ ____
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Plot by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY level. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a plot for each value of Zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as 01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

3.1.2 - UNIX Single System Inquiries

The MW OS Agent Analyzer provides the following inquiries that report on data from single UNIX systems:

1 - UNIX Global Bottlenecks (AUMCS4)
2 - UNIX Global Disk Summary (AUMCS6)

3.1.2.1 - UNIX Global Bottlenecks (AUMCS4)

The UNIX Global Bottlenecks plot uses measurements from the CA MICS AUMGPR file to graphically display an overview of processor, memory, and peak disk utilization. These metrics are displayed by system over the date and/or time intervals you request.

Each color plot displays peak disk (GPRMPBDU), CPU (GPRPCCPU), swap space (GPRPCSWS), and memory (GPRPCMEM) utilization as line values over time. The left vertical axis represents utilization percentages. Memory management I/O activity rates (GPRDMMRT) are portrayed by shaded areas on the plots. The right vertical axis shows the I/O rate in seconds.

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMCS4, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result will be displayed on your terminal.

```
---------------------- UNIX Global Bottlenecks -----------------------
 Command ===>

 Composing CA MICS Inquiry:

 Database ===> ____________________________________________
 Cycle(s) ===> __ - __
 Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
 Plot by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
 Generate a plot for each value of Zone     ===> _ (Y/N)

 Data Selection:
                         Date Range 1         Date Range 2
                     dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
   Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
   Hour         ===> __ - __    __ - __    __ - __
   Zone         ===> _ _ _ _ _ _ _ _ _ _
   Sysid        ===> ____ ____ ____ ____ ____ ____
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Plot by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY levels. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a plot for each value of Zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as 01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

3.1.2.2 - UNIX Global Disk Summary (AUMCS6)

The UNIX Global Disk Summary uses measurements from the CA MICS AUMGPR file to display an overall view of I/O activity and utilization at a system level over the date and/or time ranges you select. The color graphic plots are displayed by system, zone, and/or date intervals based on the report summarization options you specify.

Each plot displays the average I/O rate on the left axis and the total disk utilization percentage on the right axis over the date/time periods selected. The types of I/O activity (physical, logical, and memory management) contributing to the I/O rate are indicated by color. Utilization percentages are shown by lines.

The following MW OS Agent Analyzer data elements are displayed:

- o GPRDKLRT (logical disk I/O rate)
- o GPRDMMRT (virtual memory management I/O rate)
- o GPRDSKRT (physical disk I/O rate)
- o GPRMPBDU (percent utilization of busiest disk)

The following data element is not available from Solaris or AIX systems: GPRDKLRT (logical disk I/O rate)

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMCS6, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result displays on your terminal.

```
---------------------- UNIX Global Disk Summary -----------------------
Command ===>

Composing CA MICS Inquiry:

Database ===> ____________________________________________
Cycle(s) ===> __ - __
Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
Plot by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
Generate a plot for each value of Zone     ===> _ (Y/N)

Data Selection:
                        Date Range 1         Date Range 2
                     dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
  Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
  Hour         ===> __ - __    __ - __    __ - __
  Zone         ===> _ _ _ _ _ _ _ _ _ _
  Sysid        ===> ____ ____ ____ ____ ____ ____
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Plot by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY levels. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a plot for each value of zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as 01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher

than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

3.1.3 - UNIX Application Inquiries

The MW OS Agent Analyzer provides the following inquiry that reports on data about UNIX applications:

1 - UNIX Application Performance Summary (AUMLS2)

3.1.3.1 - UNIX Application Performance Summary (AUMLS2)

The UNIX Application Performance Summary uses measurements from the CA MICS AUMAPP file to display an overall view of application performance in terms of throughput and utilization by system for the date and/or time intervals you request.

Each report provides metrics listed by application name. A separate report page is produced for each logical combination of system, zone, and date/time interval, depending on the report options you specify. Measurements that allow an overall view of application load, memory utilization, I/O activity, and throughput are provided. Wait statistics are provided when you specify that the report be produced for the DETAIL timespan.

The following MW OS Agent Analyzer data elements display:

- AUMAPU (application name)
- AUMAPNUM (application number)
- APPAVAPR (average processes active)
- APPAVPRO (average processes alive)
- APPDSKIO (physical I/Os)
- APPLIOCT (logical I/Os)
- APPPCCPU (total CPU use)
- APPPCMPR (main memory use - private resident)
- APPPCMRS (main memory use - resident)
- APPPCSCP (% CPU wait time)
- APPPCSDK (% disk I/O wait time)
- APPPCSDR (% time directed to wait)
- APPPCSIM (% software impede wait time)
- APPPCSIP (% interprocess communication wait time)
- APPPCSLN (% LAN I/O wait time)
- APPPCSOT (% non-LAN, disk, terminal I/O wait time)
- APPPCSWP (% memory wait time)
- APPPCSYS (% CPU use for system processing)
- APPPCUSR (% CPU use in user mode)
- APPPKBRT (I/O transfer rate)
- DURATION (duration)

For HP-UX B.10 systems, APPMEMPR (main memory use - private resident) is not available.

For HP-UX A.09 systems, APPPKBCT (I/O transfer rate) and APPMEMRS (memory use - resident) are not available.

For Solaris and AIX systems, the following data elements are not available:

- APPPCMPR (main memory use - private resident)
- APPPCMRS (main memory use - resident)
- APPPCSCP (% CPU wait time)
- APPPCSDK (% disk I/O wait time)
- APPPCSDR (% time directed to wait)
- APPPCSIM (% software impede wait time)
- APPPCSIP (% interprocess communication wait time)
- APPPCSLN (% LAN I/O wait time)
- APPPCSOT (% non-LAN, disk, terminal I/O wait time)
- APPPCSWP (% memory wait time)

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMLS2, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result appears on your terminal and is illustrated at the end of this section.

Note: For this report, you can control the width of the report page (for example, to choose landscape versus portrait output) using the MICF Foreground Execution Parameters (MWF 2;0;2). On the line
Subset Display On ===>
type Page. Fill in a Page Width value that suits your reporting needs (80 to produce an 80 column report, for instance).

```
---------------- UNIX Application Performance Summary ---------------
Command ===>

Composing CA MICS Inquiry:

Database ===> ____________________________________________
Cycle(s) ===> __ - __
Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
Report by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
Generate a report for each value of Zone     ===> _ (Y/N)

Data Selection:
                        Date Range 1         Date Range 2
                     dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
  Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
  Hour         ===> __ - __    __ - __    __ - __
  Zone         ===> _ _ _ _ _ _ _ _ _ _
  Sysid        ===> ____ ____ ____ ____ ____ ____

  Application
  Name(AUMAPU) ===> ____________________   ____________________
                    ____________________   ____________________
                    ____________________   ____________________
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Report by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY levels. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a report for each value of Zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date

ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as 01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

Application Name (AUMAPU) - Optional. No default. 1-20 character application name (AUMAPU). You can specify up to six application names. An input observation is selected if its application name value is equal to any of the specified values.

Detail Unix Application Performance Summary 09:46 Monday, June 23, 2008
CA MICS Data Integration Applications - a2
Reporting Period: 25SEP07:00:00 - 19MAY08:23:59

SYSID=CS00 OS=HP-UX DATE=02-10-14
(continued)

Time	Application Name	Application Number	Duration	% CPU Use Total	% CPU Use System Mode	% CPU Use User Mode	% Memory Use Private Resident	% Memory Use Resident
0:20	peoplesoft	2	0:05:00	0.04	0.02	0.02	.	12.24
0:25	network	4	0:05:00	0.01	0.01	0.00	.	0.26
0:25	other	1	0:05:00	0.45	0.20	0.25	.	3.63
0:25	other_user_root	6	0:05:00	4.97	2.82	2.15	.	2.22
0:25	peoplesoft	2	0:05:00	0.04	0.02	0.02	.	12.24
0:30	network	4	0:05:00	0.01	0.01	0.00	.	0.26
0:30	other	1	0:05:00	0.34	0.14	0.20	.	3.62
0:30	other_user_root	6	0:05:00	4.99	2.85	2.14	.	2.22
0:30	peoplesoft	2	0:05:00	0.04	0.02	0.02	.	12.24
0:35	network	4	0:05:00	0.01	0.01	0.00	.	0.26
0:35	other	1	0:05:00	0.30	0.10	0.20	.	3.63

Time	Application Name	Logical I/Os	Physical I/Os	I/O Transfer Rate (KBytes)	Avg Proc Alive	Avg Proc Active	% Wait Time CPU	% Wait Time Disk I/Os	% Wait Time Memory
0:20	peoplesoft	0	0	0.00	20	9.80	0.00	0.00	0.00
0:25	network	0	0	0.00	27	5.00	0.00	0.00	0.00
0:25	other	9000	420	0.00	41	3.30	0.00	0.00	0.00
0:25	other_user_root	306300	1830	0.00	85	26.90	0.01	0.00	0.00
0:25	peoplesoft	0	0	0.00	20	9.60	0.00	0.00	0.00
0:30	network	0	60	0.00	27	5.40	0.00	0.00	0.00
0:30	other	6900	420	0.00	41	3.70	0.00	0.00	0.00
0:30	other_user_root	305700	1770	0.00	85	26.40	0.02	0.00	0.00
0:30	peoplesoft	0	0	0.00	20	9.60	0.00	0.00	0.00
0:35	network	0	0	0.00	27	4.70	0.00	0.00	0.00
0:35	other	5400	390	0.00	40	3.40	0.06	0.00	0.00

Time	Application Name	% Wait Time Software Impedes	% Wait Time Interproc Comm	% Wait Time LAN I/Os	% Wait Time Directed	% Wait Time Other I/Os
0:20	peoplesoft	0.00	2.39	0.00	4.10	0.00
0:25	network	0.00	0.00	0.00	66.66	0.00
0:25	other	11.62	2.97	0.00	21.71	0.00
0:25	other_user_root	0.36	0.74	0.00	28.42	0.19
0:25	peoplesoft	0.00	2.39	0.00	4.10	0.00
0:30	network	0.00	0.00	0.00	66.66	0.00
0:30	other	11.62	2.97	0.00	21.71	0.00
0:30	other_user_root	0.37	0.75	0.00	28.96	0.19
0:30	peoplesoft	0.00	2.39	0.00	4.10	0.00
0:35	network	0.00	0.00	0.00	66.66	0.00
0:35	other	11.61	2.98	0.00	22.47	0.01

3.1.4 - UNIX Process Activity Analysis Inquiries

The MW OS Agent Analyzer provides the following inquiry that reports on data about UNIX processes:

1 - UNIX Detail Killed Process (AUMLS3)

3.1.4.1 - UNIX Detail Killed Process (AUMLS3)

The UNIX Detail Killed Process report uses measurements from the AUMKPS file to display a detailed view of killed process activities in terms of resource consumption and performance by system for the date and/or time intervals you request.

Each report provides metrics listed by program name and measurement timestamp. A separate report page is produced for each logical combination of system and zone, depending on the report options you specify. Measurements that allow an overall view of process load, I/O activity, and wait statistics on various system resources are provided.

For HP-UX B.10 systems, the following MW OS Agent Analyzer data elements display:

- AUMAPU (application name)
- AUMLNTTY (logon TTY)
- AUMPIN (PIN)
- AUMPRI (priority)
- AUMPROG (program name)
- AUMUSER (user)
- DATE (date)
- KPSCPUTM (total CPU time)
- KPSDSKIO (physical disk I/Os)
- KPSLDKCT (logical disk I/Os)
- KPSRUNTM (run time)
- KPSSCPTM (CPU wait time)
- KPSSDKTM (disk I/O wait time)
- KPSSDRTM (directed wait time)
- KPSSIMTM (software impedes wait time)
- KPSSIPTM (IPC wait time)
- KPSSLNTM (LAN I/O wait time)
- KPSSNFTM (NFS wait time)
- KPSSOITM (other I/O wait time)
- KPSSSYTM (system wait time)
- KPSSVMTM (virtual memory wait time)
- SYSID (system ID)
- TIME (end time)

For Solaris and AIX systems, the following data elements are not available:

- KPSDSKIO (physical disk I/Os)
- KPSLDKCT (logical disk I/Os)
- KPSSCPTM (CPU wait time)
- KPSSDKTM (disk I/O wait time)
- KPSSDRTM (directed wait time)
- KPSSIMTM (software impedes wait time)
- KPSSIPTM (IPC wait time)
- KPSSLNTM (LAN I/O wait time)

- KPSSNFTM (NFS wait time)
- KPSSOITM (other I/O wait time)
- KPSSSYTM (system wait time)
- KPSSVMTM (virtual memory wait time)

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMLS3, whose Catalog Group is Perform
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result appears on your terminal and is illustrated at the end of this section.

Note: For this report, you can control the width of the report page (for example, to choose landscape versus portrait output) using the MICF Foreground Execution Parameters (MWF 2;0;2). On the line

```
  Subset Display On ===>
```

type Page. Fill in a Page Width value that suits your reporting needs (80 to produce an 80 column report, for instance).

```
------------------  UNIX Detail Killed Process Report  ---------------
 Command ===>

 Composing CA MICS Inquiry:

 Database ===> ____________________________________________
 Cycle(s) ===> __ - __
 Timespan ===> DETAIL
 Generate a report for each value of Zone     ===> _ (Y/N)

 Data Selection:          Date Range 1         Date Range 2
                       dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
   Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
   Hour         ===> __ - __    __ - __    __ - __
   Zone         ===> _ _ _ _ _ _ _ _ _
   Sysid        ===> ____ ____ ____ ____ ____ ____

   Application  ===> ____________________   ____________________
   Name(AUMAPU)      ____________________   ____________________

   User Name    ===> _________________   _________________   ______________
   Program Name ===> _________________   _________________   ______________
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

The timespan for this report is fixed as the DETAIL timespan.

Generate a report for each value of Zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as 01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value

is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

Application Name (AUMAPU) - Optional. No default. 1-20 character application name (AUMAPU). You can specify up to four application names. An input observation is selected if its application name value is equal to any of the specified values.

User Name - Optional. No default. 1-16 character user logon name (AUMUSER). You can specify up to three user logon names. An input observation is selected if its user name value is equal to any of the specified values.

Program Name - Optional. No default. 1-16 character process program name (AUMPROG). You can specify up to three process program names. An input observation is selected if its process program name value is equal to any of the specified values.

CA MICS Data Integration Applications
Detail Unix Detail Killed Process Report 10:23 Monday, June 23, 2008

SYSID=CS00 OS=HP-UX DATE=02-10-14

Time Ended	Program Name	User	PIN	Priority	Logon TTY	Application Name	Run Time	Total CPU Time
21:06	vbda	root	27898	179	?	other_user_root	0:01:42.44	0:00:03.45
21:06	vbda	root	28342	178	?	other_user_root	0:00:10.10	0:00:00.26
21:06	vbda	root	27206	178	?	other_user_root	0:03:51.05	0:00:04.87
21:06	bdf	psoft	28809	178	?	peoplesoft	0:00:37.57	0:00:00.07
21:09	runmqchl	mqm	29763	137	?	other	0:00:01.79	0:00:00.35
21:13	vbda	root	28563	179	?	other_user_root	0:06:46.52	0:00:06.68
21:16	vbda	root	2574	178	?	other_user_root	0:00:28.93	0:00:00.75
21:16	vbda	root	28918	178	?	other_user_root	0:08:47.56	0:00:07.89
21:16	vbda	root	2946	178	?	other_user_root	0:00:16.68	0:00:00.51
21:19	vbda	root	28549	178	?	other_user_root	0:14:14.91	0:00:17.29

Time Ended	Program Name	Physical Disk I/Os	Logical Disk I/Os	Wait Time CPU	Wait Time Directed	Wait Time Disk I/O	Wait Time LAN I/O	Wait Ti Other I
21:06	vbda	8471	5895	0:00:00.28	0:01:21.64	0:00:08.54	0:00:00.00	0:00:08.
21:06	vbda	690	751	0:00:00.00	0:00:06.51	0:00:01.15	0:00:00.00	0:00:02.
21:06	vbda	10976	7346	0:00:00.53	0:03:15.51	0:00:16.01	0:00:00.00	0:00:14.
21:06	bdf	579	15	0:00:00.00	0:00:01.43	0:00:16.19	0:00:00.00	0:00:19.
21:09	runmqchl	319	325	0:00:00.00	0:00:00.93	0:00:00.08	0:00:00.00	0:00:00.
21:13	vbda	19336	9394	0:00:00.65	0:04:59.20	0:00:14.39	0:00:00.00	0:01:25.
21:16	vbda	1888	1312	0:00:00.01	0:00:16.52	0:00:05.69	0:00:00.00	0:00:05.
21:16	vbda	16017	11000	0:00:00.53	0:08:14.32	0:00:08.81	0:00:00.00	0:00:15.
21:16	vbda	1329	956	0:00:00.03	0:00:11.32	0:00:01.48	0:00:00.00	0:00:03.
21:19	vbda	52319	27650	0:00:01.11	0:09:47.66	0:01:35.41	0:00:00.00	0:02:33.

Time Ended	Program Name	Wait Time SW Impedes	Wait Time Virt Memory	Wait Time NFS	Wait Time IP Comm	Wait Time System
21:06	vbda	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00
21:06	vbda	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.03
21:06	vbda	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00
21:06	bdf	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.07
21:09	runmqchl	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.08
21:13	vbda	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00
21:16	vbda	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00
21:16	vbda	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00
21:16	vbda	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00
21:19	vbda	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.00	0:00:00.08

INQUIRY: AUMLS3 RUN DATE: 23JUN08

3.1.5 - UNIX Disk Inquiries

The MW OS Agent Analyzer provides the following inquiries that report on data about disk applications:

1 - UNIX Top 5 Disk Utilization Comparison (AUMCSB)
2 - UNIX Disk I/O Activity (AUMCSC)
3 - UNIX Logical Volume Activity (AUMCSD)

3.1.5.1 - UNIX Top 5 Disk Utilization Comparison (AUMCSB)

The UNIX Top 5 Disk Utilization Comparison uses data from the CA MICS AUMDDA file to display disk utilization for the top 5 devices over the date and/or time ranges selected. The color graphic plots are displayed by system, zone, and/or date intervals based on the report summarization options you specify.

Each plot displays the average disk utilization percentage (DDAPCDUT) on the left axis over the date/time periods selected. The top five devices are those with the highest average utilization in the entire range of selected data. Utilization percentages for each device are represented by colored lines.

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMCSB, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result will be displayed on your terminal.

```
---------------- UNIX Top 5 Disk Utilization Comparison -------------
Command ===>

Composing CA MICS Inquiry:

Database ===> ____________________________________________
Cycle(s) ===> __ - __
Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
Plot by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
Generate a plot for each value of Zone     ===> _ (Y/N)

Data Selection:
                        Date Range 1         Date Range 2
                     dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
  Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
  Hour         ===> __ - __    __ - __    __ - __
  Zone         ===> _ _ _ _ _ _ _ _ _ _
  Sysid        ===> ____ ____ ____ ____ ____ ____
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Plot by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY level. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a plot for each value of zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing very often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as 01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last

value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

3.1.5.2 - UNIX Disk I/O Activity (AUMCSC)

The UNIX Disk I/O Activity chart uses data from the CA MICS AUMDDA file to display disk activity rates by I/O type for selected disk devices over the date and/or time ranges you select. The color graphic charts are displayed by system, zone, and/or date intervals based on the report summarization options you specify.

Each chart displays logical and physical disk I/O activity rates as bars. The number of I/Os per second is shown on the left axis. Physical and logical I/O rates are displayed over the date/time periods you select. The type of physical I/O--file, memory management, system, and raw--are delineated by color.

This report is not supported for Solaris or AIX systems because the elements needed are not available from the data source.

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMCSC, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result displays on your terminal

```
---------------------- UNIX Disk I/O Activity ------------------------
 Command ===>

 Composing CA MICS Inquiry:

 Database ===> ____________________________________________
 Cycle(s) ===> __ - __
 Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
 Plot by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
 Generate a plot for each value of Zone     ===> _ (Y/N)

 Data Selection:
                         Date Range 1         Date Range 2
                      dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
   Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
   Hour         ===> __ - __    __ - __    __ - __
   Zone         ===> _ _ _ _ _ _ _ _ _ _
   Sysid        ===> ____ ____ ____ ____ ____ ____

   Device
   Name         ===> _________ __________ _________
                     _________ __________ _________
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Plot by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY level. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a plot for each value of Zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as

01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

Disk Name - Optional. No default. 1-20 character disk name (DDANAME). You can specify up to six disk names. An input observation is selected if its disk name value is equal to any of the specified values.

3.1.5.3 - UNIX Logical Volume Activity (AUMCSD)

The UNIX Logical Volume Activity chart uses data from the CA MICS AUMVLM file to display logical volume activity (disk space used and logical read/write rates) for selected logical volumes over the date and/or time ranges you select.

Each color plot displays the logical read rate (VLMPSLRD) and logical write rate (VLMPSLWT) as line values over time. The left vertical axis represents the I/O rate per second. Disk space used (VLMPCSPU) is portrayed by a shaded area on the plots. The right vertical axis shows the utilization percentage.

This report is not supported for Solaris or AIX systems because the elements needed are not available from the data source.

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMCSD, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result displays on your terminal.

```
---------------------- UNIX Logical Volume Activity --------------------
 Command ===>

 Composing CA MICS Inquiry:

 Database ===> ____________________________________________
 Cycle(s) ===> __ - __
 Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
 Plot by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
 Generate a plot for each value of Zone     ===> _ (Y/N)

 Data Selection:
                         Date Range 1         Date Range 2
                      dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
   Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
   Hour         ===> __ - __    __ - __    __ - __
   Zone         ===> _ _ _ _ _ _ _ _ _ _
   Sysid        ===> ____ ____ ____ ____ ____ ____

   Logical
   Volume Name  ===> ____________________ ____________________
                     ____________________ ____________________
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Plot by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY level. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a plot for each value of Zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as

01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

Logical Volume Name - Optional. No default. 1-20 character logical volume name (VLMNAME). You can specify up to four logical volumes. An input observation is selected if its logical volume name value is equal to any of the specified values.

3.1.6 - UNIX Global Inquiries

The MW OS Agent Analyzer provides the following inquiry that reports on global data:

1 - UNIX System Performance Summary (AUMLS1)

3.1.6.l - UNIX System Performance Summary (AUMLS1)

The UNIX System Performance Summary uses measurements from the CA MICS AUMGPR file to display an overall view of system utilization and workload throughput by system for the date and/or time intervals you request.

Each report is organized into categories of related measurements that allow an overall view of system utilization, load, and performance. Categories include sections for processor load, memory utilization, I/O activity, and workload throughput.

For HP-UX systems, the following MW OS Agent Analyzer data elements are displayed:

- GPRAVAPR (active processes)
- GPRAVCPQ (run queue length)
- GPRAVDKQ (disk queue length)
- GPRAVIPQ (IPC resources queue length)
- GPRAVMMQ (memory queue length)
- GPRAVNTQ (network queue length)
- GPRAVPRO (alive processes)
- GPRAVPRQ (priority queue length)
- GPRAVUSR (user sessions)
- GPRDKLRT (logical I/O)
- GPRDMMRT (virtual memory management I/O rate)
- GPRDMRRT (memory management - reads)
- GPRDMWRT (memory management - writes)
- GPRDSKRT (physical I/O)
- GPRPAGRT (page fault rate)
- GPRPCCPU (total CPU busy)
- GPRPCCSW (CPU use - context switching)
- GPRPCHIT (% logical reads met by cache)
- GPRPCINT (CPU use - interrupt handling)
- GPRPCMEM (main memory % used)
- GPRPCMMS (main memory % used - system)
- GPRPCMMU (main memory % used - users)
- GPRPCNIC (CPU use - nice priority)
- GPRPCNOR (CPU use - normal priority)
- GPRPCRTM (CPU use - real time priority)
- GPRPCSSC (CPU use - system calls)
- GPRPCSWS (disk swap space % used)
- GPRPCSYS (CPU use - system mode)
- GPRPCUSR (CPU use - user mode)

For Solaris and AIX systems, the following data elements are displayed:

- GPRAVAPR (active processes)
- GPRAVCPQ (run queue length)
- GPRAVUSR (user sessions)
- GPRDMMRT (virtual memory management I/O rate)
- GPRDSKRT (physical I/O)
- GPRPAGRT (page fault rate)
- GPRPCCPU (total CPU busy)
- GPRPCMEM (main memory % used)
- GPRPCMMS (main memory % used - system)
- GPRPCMMU (main memory % used - users)
- GPRPCSWS (disk swap space % used)
- GPRPCSYS (CPU use - system mode)
- GPRPCUSR (CPU use - user mode)

To produce the report:

- o Enter the CA MICS Workstation Facility (MWF)
- o Choose option 2, CA MICS Information Center Facility (MICF)
- o Choose option 2, Database Inquiries
- o Select AUMLS1, whose Catalog Group is Perform
- o Complete the MICF Execution panel
- o Complete the specification panel illustrated below (Descriptions of the fields follow the panel)

The result will be displayed on your terminal and is illustrated at the end of this section.

Note: For this report, the page width must be 132. Examine, and if necessary, change the page width using the MICF Foreground Execution Parameters (MWF 2;0;2). On the line

```
  Subset Display On ===>
```

type Page. For Page Width, ensure that the value is set to 132.

```
---------------------- UNIX System Performance Summary ---------------
Command ===>

Composing CA MICS Inquiry:

Database ===> ____________________________________________
Cycle(s) ===> __ - __
Timespan ===> ______ (DETAIL/DAYS/WEEKS/MONTHS)
Report by hour if timespan is DAYS or WEEKS ===> _ (Y/N)
Generate a report for each value of Zone     ===> _ (Y/N)

Data Selection:
                        Date Range 1         Date Range 2
                     dd mmm yy   dd mmm yy   dd mmm yy   dd mmm yy
  Date         ===> __ ___ __ - __ ___ __   __ ___ __ - __ ___ __
  Hour         ===> __ - __    __ - __    __ - __
  Zone         ===> _ _ _ _ _ _ _ _ _ _
  Sysid        ===> ____ ____ ____ ____ ____ ____
```

Data base - Required. No default. Specify up to six data base unit IDs (the IDs of the CA MICS data base units from which files are read). Duplicates are not allowed. If CA MICS files exist in only one data base unit, this parameter is set automatically. To display a selection list of valid data base unit IDs, specify ? (question mark).

Cycle(s) - Required. Defaults to 01. This is the cycle or cycles of the specified time-span of the CA MICS files that are used as input to the analysis. Two data entry fields are provided so you can specify a range of cycles. The fields can contain any numeric value between 0 and 99. If you want only a single cycle (such as cycle 01), specify the selected cycle in the first entry field and leave the second field blank. Leading zeros are not required and are inserted automatically if you supply a single-digit number (7 becomes 07). If you specify a cycle range, you may specify it in either ascending or descending order (for example, 01-05 or 05-01).

Time-span - Required. Default varies by report time-span. The time-span of the CA MICS data base from which the primary input files are selected. Required input file(s) must be available at the specified time-span.

Report by hour - Required. Defaults to Y. At the DAYS and WEEKS timespans, reports can be produced at an hourly level or at the DAILY or WEEKLY levels. If you specify Y (Yes) to Report by hour, the data is summarized and displayed by hour and separate reports are produced for each day or week. If you specify N (No) to Report by hour, the data is summarized by day or week and a single report listing all days or weeks is produced.

Generate a report for each value of Zone - Required. Defaults to Y. Specify Y (Yes) to generate reports for each zone. The nature and resource requirements of UNIX processing often change from one zone to another (for example, prime shift online demands versus production batch requirements at non-prime shift periods). Zone reporting gives you the ability to display separate reports for each observed zone. If you specify N (No), the data is summarized to the appropriate level and displayed on a single report.

Date - Optional. No default. Up to two date ranges in DD MMM YY format. The first value in each date range is the starting date and the last value is the ending date. The ending date must be higher than the starting date. The date ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid date range specification is specifying date range 1 as 01 JAN 95 - 31 DEC 95 and specifying date range 2 as 01 JAN 94 - 31 DEC 94. An input observation is selected if its date value is greater than or equal to the starting date and is less than or equal to the ending date of any specified date range.

Hour - Optional. No default. 0-23. Up to three hour

ranges. The first value is the starting hour and the last value is the ending hour. The ending hour must be higher than the starting hour. The hour ranges should not overlap, but you can specify the higher range before the lower range. For example, a valid hour range specification is 18 - 23 and 09 - 17. An input observation is selected if its hour value is greater than or equal to the starting hour and is less than or equal to the ending hour of any specified hour range.

Zone - Optional. No default. 1-character zone value from 1 to 9. You can specify up to nine zone values. An input observation is selected if its zone value is equal to any of the specified values.

Sysid - Optional. No default. 1-4 character SYSID value. You can specify up to six SYSIDs. An input observation is selected if its SYSID value is equal to any of the specified values.

Detail Unix System Performance Summary
CA MICS Data Integration Applications

SYSID=CS00 OS=HP-UX

INTVL-DAY (HR:MM-DD):	0:59-14	1:05-14	1:09-14	1:15-14	1:19-14	1:25-14	1:30-14	1:35-14	1:40-14	1:44-14	CHANGE %
CPU UTILIZATION PERCENTS											
System Mode	3.95	7.12	4.06	4.06	3.98	3.97	3.98	6.23	3.97	6.20	56.96 %
-- System Calls	2.99	5.94	3.07	3.05	3.01	2.98	2.99	5.16	2.99	5.16	72.58 %
-- Context Switching	0.10	0.14	0.10	0.09	0.09	0.09	0.09	0.12	0.09	0.12	20.00 %
-- Interrupt Handling	0.82	0.96	0.85	0.84	0.84	0.82	0.84	0.89	0.84	0.88	7.32 %
User Mode	2.42	3.95	2.45	2.42	2.44	2.42	2.46	3.66	2.47	3.68	52.07 %
-- Normal Priority	1.99	3.45	2.02	2.01	2.01	2.00	2.02	3.20	2.03	3.21	61.31 %
-- Nice Priority	0.14	0.14	0.14	0.13	0.14	0.13	0.15	0.12	0.15	0.13	-7.14 %
-- Real Time Priority	0.29	0.36	0.29	0.29	0.29	0.29	0.29	0.35	0.29	0.34	17.24 %
Total CPU Busy	6.37	11.07	6.51	6.48	6.42	6.39	6.44	9.89	6.44	9.88	55.10 %
MEMORY UTILIZATION											
Main Memory % Used	35.07	35.09	35.09	35.09	35.07	35.09	35.09	35.09	35.07	35.09	0.06 %
-- Users	18.66	18.68	18.68	18.68	18.66	18.68	18.68	18.68	18.66	18.68	0.11 %
-- System	16.41	16.41	16.41	16.41	16.41	16.41	16.41	16.41	16.41	16.41	---
Disk Swap Space % Used	23.00	23.00	23.00	23.00	23.00	23.00	23.00	23.00	23.00	23.00	---
Page Fault Rate	163.30	386.60	165.90	168.10	168.40	171.00	167.70	362.00	165.50	356.00	118.00 %
WORKLOAD AVERAGES											
User Sessions	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
Alive Processes	175.00	175.00	175.00	175.00	175.00	175.00	175.00	175.00	175.00	175.00	---
Active Processes	64.00	68.00	64.00	65.00	66.00	66.00	64.00	65.00	66.00	67.00	4.69 %
I/O LOAD PER SECOND											
Logical I/O	1046.00	1252.00	1048.00	1049.00	1047.00	1050.00	1048.00	1230.00	1049.00	1234.00	17.97 %
Physical I/O	10.00	65.30	12.50	10.40	10.50	9.90	10.20	24.60	10.80	23.30	133.00 %
Memory Management I/O	0.00	0.20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	N/A
Cache Hit Ratio	99.98	98.76	99.80	100.00	99.97	99.99	100.00	99.88	99.99	99.97	-0.01 %

Detail Unix System Performance Summary
CA MICS Data Integration Applications

SYSID=CS00 OS=HP-UX

INTVL-DAY (HR:MM-DD):	0:59-14	1:05-14	1:09-14	1:15-14	1:19-14	1:25-14	1:30-14	1:35-14	1:40-14	1:44-14 CH
AVERAGE QUEUE LENGTH										
Run	0.06	0.05	0.10	0.09	0.12	0.10	0.04	0.04	0.16	0.16
Priority	0.00	0.02	0.00	0.00	0.02	0.00	0.01	0.01	0.01	0.02
Disk	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Memory	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
IPC Resources	2.90	2.85	3.08	2.88	2.85	2.87	2.85	2.85	2.88	2.85
Network	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

On this report, a Change % of N/A means that a percentage cannot be calculated because the divisor is 0.

A Change % of --- means that the percentage is 0 or that there was no data to compare.

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 4

RESERVED

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 4

RESERVED

Chapter 4. RESERVED

The CA MICS Analyzer for MeasureWare does not distribute exceptions.

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 5

FILES

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 5

FILES

Chapter 5. FILES

The Analyzer stores its data in the UNIX information area, whose three-character information area identifier is AUM.

The table in Figure 5-1 lists the product's files and the timespans that are activated for each file. For each file in the information area, the following information is provided:

XDWMYT - Defines the timespans in which the file is supported.

X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported

File - The SAS name used to access this file.

File Name - The descriptive label for the file.

MeasureWare Information Area	Date Generated: Thu, Jul 18, 2019

XDWMYT	File	File Name
XDWMY.	AUMAPP	UNIX APPLICATION PERFORMANCE FILE
XDWMY.	AUMDDA	UNIX DISK ACTIVITY FILE
XDWMY.	AUMGPR	UNIX GLOBAL PERFORMANCE FILE
X.....	AUMIPS	UNIX PROCESS EXCEPTION FILE
X..MY.	AUMKPS	UNIX PROCESS ACTIVITY FILE
X..M..	AUMMCF	UNIX SYSTEM AND COLLECTOR PROFILE
XDWMY.	AUMNIF	UNIX LAN ACTIVITY FILE
.....	AUMTTA	UNIX BUSINESS TRANSACTION SLO ACTIVITY
XDWMY.	AUMVLM	UNIX LOGICAL VOLUME ACTIVITY FILE

Figure 5-1. UNIX Information Area Files

This chapter contains the following sections:

1 - Data Element Naming Conventions
2 - UNIX Global Performance File (AUMGPR)
3 - UNIX Application Performance File (AUMAPP)
4 - UNIX Process Activity File (AUMKPS)
5 - UNIX Process Exception File (AUMIPS)
6 - UNIX LAN Activity File (AUMNIF)
7 - UNIX Disk Activity File (AUMDDA)
8 - UNIX System and Collector Profile (AUMMCF)
9 - UNIX Logical Volume File (AUMVLM)
10 - UNIX Business Transaction SLO Activity File (AUMTTA)

5.1 - Data Element Naming Conventions

There are two types of CA MICS data elements: standard and common. Standard data elements are unique to a CA MICS file, while common data elements are used by more than one CA MICS file.

The first three characters of a standard data element's name are the three-character file identifier to which the element belongs.

Common data elements do not use a data element prefix. They have a common definition across database information areas or across files within an information area. Common data elements are listed in the Sequence/Summary Data Elements and Common Data Elements sections of the data elements lists that accompany the following file documentation.

The following chart lists each file's descriptive name, its six-character file name, and the three-character standard data element name prefix.

Descriptive Name	File Name	Element Prefix
UNIX Global Performance File	AUMGPR	GPR
UNIX Application Performance File	AUMAPP	APP
UNIX Process Activity File	AUMKPS	KPS
UNIX Process Exception File	AUMIPS	IPS
UNIX LAN Activity File	AUMNIF	NIF
UNIX Disk Activity File	AUMDDA	DDA
UNIX System and Collector Profile	AUMMCF	MCF
UNIX Logical Volume Activity File	AUMVLM	VLM
UNIX Business Transaction SLO Activity	AUMTTA	TTA

5.2 - UNIX Global Performance File (AUMGPR)

The UNIX Global Performance file contains global system measurements useful for capacity, system load, and performance analyses. It includes information about CPU use, disk I/Os, paging and swapping, and wait queues, and is derived from MeasureWare Global metrics.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMGPR File Organization
2 - AUMGPR Data Elements List
3 - AUMGPR Usage Considerations
4 - AUMGPR Retrieval Examples

5.2.1 - AUMGPR File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	YEAR	MONTH	DAY
	HOUR	ENDTS			
DAYS	SYSID	AUMOPSYS	YEAR	MONTH	DAY
	HOUR				
WEEKS	SYSID	AUMOPSYS	YEAR	WEEK	ZONE
	HOUR				
MONTHS	SYSID	AUMOPSYS	YEAR	MONTH	ZONE
YEARS	SYSID	AUMOPSYS	YEAR	ZONE	
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.
NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-2. AUMGPR Timespan Granularity Chart

5.2.2 - AUMGPR Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
Sequence/Summary Data Elements		
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	SYSID	- System Identifier
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone

Common Data Elements

```
XDWMY.E  AUMCPMDL - System Model Identification
XDWMY.E  AUMSAMPL - Number of Data Samples
XDWMY.E  AUMSYSID - Name of Host System
XD....E  DAYNAME  - Name of Day of Week
XDWMY.E  DURATION - Recording Interval Time
XDWMY.E  ENDTS    - End Time Stamp
XDWMY.E  INTERVLS - Number of Recording Intervals
XDWMY.E  STARTTS  - Start Time Stamp
```

Retained Data Elements

```
XDWMY.E  GPRACCPU - CPUs Online per System
XDWMY.E  GPRSUPTM - Time of Last System Reboot
```

Accumulated Data Elements

```
XDWMY.E  GPRACTVM - Active Virtual Memory Size
XDWMY.E  GPRAPRCT - Sum of Avg Processes Active
XDWMY.E  GPRBIOCT - Processes Blocked on Local Disks
XDWMY.E  GPRCACHE - Size of Buffer Cache Configured
XDWMY.E  GPRCHUSE - The Size of Currently Used Buffers
XDWMY.E  GPRCPQTM - CPU Queue Wait Time
XDWMY.E  GPRCPRCT - Total Processes Completed
XDWMY.E  GPRCPUCP - Total Available CPU Time
XDWMY.E  GPRCPUTM - Total CPU Busy
XDWMY.E  GPRCSWTM - CPU Busy for Context Switching
XDWMY.E  GPRDBKCT - Block I/Os
XDWMY.E  GPRDBRCT - Block Reads
XDWMY.E  GPRDBWCT - Block Writes
XDWMY.E  GPRDFRCT - Disk Reads by File System
XDWMY.E  GPRDFWCT - Disk Writes by File System
XDWMY.E  GPRDKQTM - Disk I/O Queue Wait Time
XDWMY.E  GPRDMMCT - Virtual Memory Management I/Os
XDWMY.E  GPRDMRCT - Disk Reads for Virt Memory Management
XDWMY.E  GPRDMWCT - Disk Writes for Virt Memory Management
XDWMY.E  GPRDRRCT - Raw Reads
XDWMY.E  GPRDRWCT - Raw Writes
XDWMY.E  GPRDSKIO - Physical Disk I/Os
XDWMY.E  GPRDSRCT - Disk Reads for File System Management
XDWMY.E  GPRDSWCT - Disk Writes for File System Management
XDWMY.E  GPRFSDCT - File System I/Os
XDWMY.E  GPRFSMCT - File System Management I/Os
XDWMY.E  GPRINODU - Inode Entries in Use
XDWMY.E  GPRINTTM - CPU Busy for Interrupt Handling
XDWMY.E  GPRIPQTM - IPC Queue Wait Time
XDWMY.E  GPRLCOLS - LAN Collisions
XDWMY.E  GPRLDKCT - Logical Disk I/Os
XDWMY.E  GPRLERRS - LAN Packet Errors
XDWMY.E  GPRLOSTB - Trace Buffers Lost
XDWMY.E  GPRLPCKI - Successful LAN Packets Received
XDWMY.E  GPRLPCKO - Successful LAN Packets Sent
XDWMY.E  GPRLRDCT - Logical Disk Reads
XDWMY.E  GPRLRDKB - KBytes Logically Read From Disk
XDWMY.E  GPRLUSED - Number of File Locks in Use
XDWMY.E  GPRLWTCT - Logical Disk Writes
XDWMY.E  GPRLWTKB - KBytes Logically Written To Disk
XDWMY.E  GPRMEMSW - Amount of Swap Space Configured
XDWMY.E  GPRMEMSZ - Amount of Main Memory Configured
XDWMY.E  GPRMEMUS - Amount of Main Memory Available for User
XDWMY.E  GPRMEMUT - Main Memory Used
XDWMY.E  GPRMMEMS - Main Memory Used by System
XDWMY.E  GPRMMEMU - Main Memory Used by Users
XDWMY.E  GPRMMQTM - Memory Queue Wait Time
XDWMY.E  GPRMQUSE - Message Queues in Use
XDWMY.E  GPRMSGQU - Message Queues Available
XDWMY.E  GPRNFSCT - Network File System Requests
XDWMY.E  GPRNICTM - CPU Busy at Nice Priorities
XDWMY.E  GPRNORTM - CPU Busy at Normal Priorities
XDWMY.E  GPRNPROC - Value of NPROC Parameter
XDWMY.E  GPRNTQTM - Network Queue Wait Time
XDWMY.E  GPROFLOW - Unmeasurable New Transactions
XDWMY.E  GPROUTQU - Outbound Queue Lengths
XDWMY.E  GPRPAUTM - CPU Idle Due to I/O Wait
XDWMY.E  GPRPFLCT - Page Faults
XDWMY.E  GPRPGOCT - Number of Page Outs
XDWMY.E  GPRPKBCT - KBytes Physically Transferred
XDWMY.E  GPRPRDCT - Physical Disk Reads
XDWMY.E  GPRPRITM - Time Blocked on Priority
XDWMY.E  GPRPROCT - Sum of Avg Processes Alive
XDWMY.E  GPRPTBLU - Proc Table Entries Used
XDWMY.E  GPRPWTCT - Physical Disk Writes
XDWMY.E  GPRRAWCT - Raw I/Os
XDWMY.E  GPRRDHIT - Buffered Read Hits
XDWMY.E  GPRRTMTM - CPU Busy at Real Time Priorities
XDWMY.E  GPRRUNTM - Total Run Time of Completing Processes
XDWMY.E  GPRSCALL - Total System Calls
XDWMY.E  GPRSEMID - Semaphore Identifiers Available
XDWMY.E  GPRSEMST - Semaphore Identifiers in Use
XDWMY.E  GPRSHMSG - Table Memory Segments Available
XDWMY.E  GPRSLQTM - Sleep Queue Wait Time
XDWMY.E  GPRSMEMS - Shared Memory Segments in Use
XDWMY.E  GPRSPRCT - Total Processes Started
XDWMY.E  GPRSSCTM - CPU Busy for System Calls
XDWMY.E  GPRSUSCT - Procs Suspended During Interval
XDWMY.E  GPRSWACT - Processes Waiting to be Swapped in
XDWMY.E  GPRSWAPU - Swap Space Used
XDWMY.E  GPRSWOCT - Number of Swap Outs/Deactivations
XDWMY.E  GPRSWPCT - Process Swaps, Activation/Deactivation
XDWMY.E  GPRSYSTM - CPU Busy in System Mode
XDWMY.E  GPRTBLEN - File Table Entries Used
XDWMY.E  GPRTBLUK - Entries Used by Kernel
```

```
 XDWMY.E  GPRUSRCT - Sum of Avg Login Users                        XDWMY.E  GPRSCLRT - System Procedure Call Rate
 XDWMY.E  GPRUSRTM - CPU Busy in User Mode                         XDWMY.E  GPRSWORT - Swap-Out Rate
                                                                   XDWMY.E  GPRSWPRT - Swapping Rate
Maximum Data Elements

 XDWMY.E  GPRMPBDU - Percent Utilization of Busiest Disk
 XDWMY.E  GPRMPDSU - Peak Disk Space Utilization
 XDWMY.E  GPRMXCPU - Max Number of Processors Configured
 XDWMY.E  GPRMXDSK - Max Number of Disk Drives Configured
 XDWMY.E  GPRMXDTM - Time Busiest Disk Was Busy
 XDWMY.E  GPRMXLAN - Max Number of LAN Interfaces Configured

Derived Data Elements

 XDWMY.E  GPRAVAPR - Avg Number of Processes Active
 XDWMY.E  GPRAVBIQ - Avg Processes Blocked on Local Disks
 XDWMY.E  GPRAVCPQ - Avg Processes in Run Queue
 XDWMY.E  GPRAVDKQ - Avg Processes Waiting for Disk I/Os
 XDWMY.E  GPRAVIPQ - Avg Processes Waiting for IPC Resources
 XDWMY.E  GPRAVLCL - Rate of LAN Collisions per Minute
 XDWMY.E  GPRAVLER - Rate of LAN Errors per Minute
 XDWMY.E  GPRAVMMQ - Avg Processes Waiting for Memory
 XDWMY.E  GPRAVNTQ - Avg Processes Waiting for Network
 XDWMY.E  GPRAVPRO - Avg Number of Processes Alive
 XDWMY.E  GPRAVPRQ - Avg Processes Blocked on Priority
 XDWMY.E  GPRAVRTM - Avg Run Time of Completing Processes
 XDWMY.E  GPRAVSLQ - Avg Processes in Sleep State
 XDWMY.E  GPRAVSUS - Avg Number of Processes Suspended
 XDWMY.E  GPRAVSWA - Avg Number of Processes Waiting Swap-In
 XDWMY.E  GPRAVUSR - Avg Number of Login Users
 XDWMY.E  GPRDKLRT - Logical Disk I/O Rate
 XDWMY.E  GPRDMMRT - Virtual Memory Management I/O Rate
 XDWMY.E  GPRDMRRT - Virt Memory Management Disk Read Rate
 XDWMY.E  GPRDMWRT - Virt Memory Management Disk Write Rate
 XDWMY.E  GPRDSKRT - Physical Disk I/O Rate
 XDWMY.E  GPRNFSRT - Rate of Network File System Requests
 XDWMY.E  GPRPAGRT - Page Fault Rate
 XDWMY.E  GPRPCAVM - Percent of Virtual Memory in Active Use
 XDWMY.E  GPRPCCPU - Percent Total CPU Busy
 XDWMY.E  GPRPCCSW - Percent CPU Busy for Context Switching
 XDWMY.E  GPRPCHIT - Cache Hit Ratio
 XDWMY.E  GPRPCINT - Percent CPU Busy for Interrupt Handling
 XDWMY.E  GPRPCMEM - Percent of Main Memory Used
 XDWMY.E  GPRPCMMS - Percent of Main Memory Used by System
 XDWMY.E  GPRPCMMU - Percent of Main Memory Used by Users
 XDWMY.E  GPRPCNIC - Percent CPU Busy at Nice Priorities
 XDWMY.E  GPRPCNOR - Percent CPU Busy at Normal Priorities
 XDWMY.E  GPRPCRTM - Percent CPU Busy at Real Time Priorities
 XDWMY.E  GPRPCSSC - Percent CPU Busy for System Calls
 XDWMY.E  GPRPCSWS - Percent of Swap Space Used
 XDWMY.E  GPRPCSYS - Percent CPU Busy in System Mode
 XDWMY.E  GPRPCUSR - Percent CPU Busy in User Mode
```

5.2.3 - AUMGPR Usage Considerations

1. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

2. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

5.2.4 - AUMGPR Retrieval Examples

Plot the use of CPU by system and users.

```
DATA GPR;
  SET DAYS.AUMGPR01;
  DATE = DATEPART(ENDTS);
  IF 8 < HOUR < 18;
  %GPRDERV;
  LABEL GPRPCUSR = '%CPU Time Charged to Users'
        GPRPCSYS = '%CPU Time System Overhead'
        GPRPCCPU = '%Total CPU Busy';
  FORMAT DATE MMDDYY8.;
RUN;

PROC TIMEPLOT DATA=GPR;
  BY SYSID DATE;
  ID HOUR;
    PLOT GPRPCCPU='T' GPRPCUSR='U'  GPRPCSYS='S'
    / OVERLAY HILOC POS=0 AXIS=0 TO 100;
    TITLE "CPU Consumption by System and Users";
  RUN;
```

5.3 - UNIX Application Performance File (AUMAPP)

The UNIX Application Performance file contains interval-based measurements useful for analyzing application use. It includes information about CPU use, disk I/Os, memory, transactions, processes, and wait states, and is derived from MeasureWare Application metrics.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMAPP File Organization
2 - AUMAPP Data Elements List
3 - AUMAPP Usage Considerations
4 - AUMAPP Retrieval Examples

5.3.1 - AUMAPP File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	AUMAPU	YEAR	MONTH
	DAY	HOUR	ENDTS		
DAYS	SYSID	AUMOPSYS	AUMAPU	YEAR	MONTH
	DAY	HOUR			
WEEKS	SYSID	AUMOPSYS	AUMAPU	YEAR	WEEK
	ZONE	HOUR			
MONTHS	SYSID	AUMOPSYS	AUMAPU	YEAR	MONTH
	ZONE				
YEARS	SYSID	AUMOPSYS	AUMAPU	YEAR	ZONE
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-3. AUMAPP Timespan Granularity Chart

5.3.2 - AUMAPP Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
XDWMY.E	AUMAPU	- Application Name
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	SYSID	- System Identifier
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
		Common Data Elements
XDWMY.E	AUMAPNUM	- Application Number
XDWMY.E	AUMSAMPL	- Number of Data Samples
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
		Retained Data Elements
X.....E	APPPCSCP	- Pct CPU Wait
X.....E	APPPCSDK	- Pct Disk I/O Wait
X.....E	APPPCSDR	- Pct Time Directed to Wait
X.....E	APPPCSIM	- Pct Software Impede Wait
X.....E	APPPCSIP	- Pct Interprocess Communication Wait
X.....E	APPPCSLN	- Pct LAN I/O Wait
X.....E	APPPCSOT	- Pct Non-LAN/DISK/TERM I/O Wait
X.....E	APPPCSWP	- Pct Memory Wait
		Accumulated Data Elements
XDWMY.E	APPAPRCT	- Sum of Avg Processes Active
XDWMY.E	APPCOST	- Processing Charges
XDWMY.E	APPCPRCT	- Number of Processes Completed
XDWMY.E	APPCPUCM	- PRM CPU Cap Mode
XDWMY.E	APPCPUCP	- Total Available CPU Time
XDWMY.E	APPCPUEN	- PRM CPU Entitlement
XDWMY.E	APPCPUST	- PRM CPU State
XDWMY.E	APPCPUTM	- Total CPU Use
XDWMY.E	APPDBKCT	- Block I/Os
XDWMY.E	APPDBRCT	- Block Reads
XDWMY.E	APPDBWCT	- Block Writes
XDWMY.E	APPDFSCT	- Number of File System Reads and Writes
XDWMY.E	APPDMMCT	- Number of Memory Manager Reads and Write
XDWMY.E	APPDRAIO	- Number of Raw Reads and Writes
XDWMY.E	APPDSKIO	- Number of Physical Disk I/Os
XDWMY.E	APPDSYIO	- Number of System Reads and Writes
XDWMY.E	APPLIOCT	- Logical Disk I/Os
XDWMY.E	APPLOGMD	- PRM Logging Mode
XDWMY.E	APPLRDCT	- Logical Disk Reads
XDWMY.E	APPLWTCT	- Logical Disk Writes
XDWMY.E	APPMEMAV	- PRM Available Memory
XDWMY.E	APPMEMEN	- PRM MEM Entitlement
XDWMY.E	APPMEMPR	- Private Resident Memory Size in KBytes
XDWMY.E	APPMEMRS	- Resident Memory Size in KBytes
XDWMY.E	APPMEMST	- PRM MEM State
XDWMY.E	APPMEMSW	- Amount of Swap Space Configured
XDWMY.E	APPMEMSZ	- Amount of Main Memory Configured

```
XDWMY.E  APPMEMUP - PRM MEM Upperbound
XDWMY.E  APPMEMUS - PRM Memory Used by Processes
XDWMY.E  APPMEMVR - Virtual Memory Size in KBytes
XDWMY.E  APPNICTM - CPU Use at Nice Priorities
XDWMY.E  APPNORTM - CPU Use at Normal Priorities
XDWMY.E  APPPFMAJ - Page Faults to Disk
XDWMY.E  APPPFMIN - Page Faults to Memory
XDWMY.E  APPPKBCT - I/O Transfer Count in KBytes
XDWMY.E  APPPRDCT - Physical Disk Reads
XDWMY.E  APPPROCT - Sum of Avg Processes Alive
XDWMY.E  APPPWTCT - Physical Disk Writes
XDWMY.E  APPRTMTM - CPU Use at Real Time Priorities
XDWMY.E  APPRUNTM - Total Run Time of Completing Processes
XDWMY.E  APPSUSCT - Procs suspended during interval
XDWMY.E  APPSYSTM - CPU Use for System Processing
XDWMY.E  APPUSRTM - CPU Use in User Mode
```

Derived Data Elements

```
XDWMY.E  APPAVAPR - Avg Processes Active
XDWMY.E  APPAVPRO - Avg Processes Alive
XDWMY.E  APPAVRTM - Avg Run Time of Completing Processes
XDWMY.E  APPAVSUS - Avg Number of Processes Suspended
XDWMY.E  APPDFSRT - Rate of File System Reads/Writes (IO/Sec
XDWMY.E  APPDMMRT - Rate of Memory Mgr. Reads/Writes (IO/Sec
XDWMY.E  APPDRART - Rate of Raw Reads/Writes (IO/Sec)
XDWMY.E  APPDSYRT - Rate of System Reads/Writes (IO/Sec)
XDWMY.E  APPMAJRT - Rate of Page Faults to Disk
XDWMY.E  APPMINRT - Rate of Page Faults to Memory
XDWMY.E  APPPCCPU - Percent Total CPU Use
XDWMY.E  APPPCMPR - Percent Memory Use for Private Resident
XDWMY.E  APPPCMRS - Percent Memory Use for Resident Memory
XDWMY.E  APPPCNIC - Percent CPU Use at Nice Priorities
XDWMY.E  APPPCNOR - Percent CPU Use at Normal Priorities
XDWMY.E  APPPCRTM - Percent CPU Use at Real Time Priorities
XDWMY.E  APPPCSYS - Percent CPU Use for System Processing
XDWMY.E  APPPCUSR - Percent CPU Use in User Mode
XDWMY.E  APPPKBRT - I/O Transfer Rate in KBytes
XDWMY.E  APPPSPRD - Physical Disk Read Rate
XDWMY.E  APPPSPWT - Physical Disk Write Rate
```

5.3.3 - AUMAPP Usage Considerations

1. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

2. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

5.3.4 - AUMAPP Retrieval Examples

Print the use of CPU, memory, and swap spaces for the top ten applications on the system.

```
%LET BY = SYSID YEAR MONTH DAY AUMAPU;
%LET BREAK = AUMAPU;
PROC SORT DATA=DAYS.AUMAPP01 OUT=APP;
  BY &BY;
  RUN;
DATA APP;
  SET DAYS.AUMAPP01;
  DATE = DATEPART(ENDTS);
  FORMAT DATE MMDDYY8.;
  SET APP;
  %APPFMT(TS=DAYS,DERV=YES);
  %APPSUM;
RUN;
%LET BY = SYSID DATE;
%LET BREAK = DATE;
PROC SORT; BY &BY DESCENDING APPCPUTM;
DATA APP2 (KEEP=&BY AUMAPU APPPCCPU APPCPUTM APPPCUSR
           APPPCNIC APPPCSYS APPPCRTM
           APPAVAPR RANK);
  SET APP;
  BY &BY;
  RETAIN RANK 0;
  IF FIRST.&BREAK THEN RANK = 0;
  RANK+1;
  IF RANK > 10 THEN RETURN;
  OUTPUT;
  LABEL RANK = 'Rank'
        AUMAPU   = 'Application*Name'
        APPPCCPU = '%CPU Use*Total'
        APPCPUTM = 'Total*CPU Use'
        APPPCUSR = '%CPU Use*User Mode'
        APPPCNIC = '%CPU Use*Nice'
        APPPCSYS = '%CPU Use*System'
        APPPCRTM = '%CPU Use*Real Time'
        APPAVAPR = 'Avg Process*Active'
        ;
RUN;
PROC PRINT SPLIT='*';
  BY &BY;
  PAGEBY &BREAK;
  ID RANK;
  VAR  AUMAPU APPCPUTM APPAVAPR APPPCCPU APPPCUSR APPPCNIC
       APPPCRTM APPPCSYS;
  TITLE "Top 10 Applications by CPU Use";
RUN;
```

5.4 - UNIX Process Activity File (AUMKPS)

The UNIX Process Activity file contains measurements about completed processes and is useful for accounting and process analysis. It includes information about CPU, disk, memory, and transactions, and is derived from MeasureWare Application metrics.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMKPS File Organization
2 - AUMKPS Data Elements List
3 - AUMKPS Usage Considerations
4 - AUMKPS Retrieval Examples

5.4.1 - AUMKPS File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID YEAR	AUMOPSYS MONTH	AUMACT1 DAY	AUMACT2 HOUR	AUMACT3 ENDTS
DAYS	N/A				
WEEKS	N/A				
MONTHS	SYSID MONTH	AUMOPSYS ZONE	AUMACT1	AUMACT2	YEAR
YEARS	SYSID ZONE	AUMOPSYS	AUMACT1	AUMACT2	YEAR
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.
NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-4. AUMKPS Timespan Granularity Chart

5.4.2 - AUMKPS Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
Sequence/Summary Data Elements		
X..MY.E	AUMACT1	- DIVISION
X..MY.E	AUMACT2	- USER GROUP
X.....E	AUMACT3	- USER NAME
X..MY.E	AUMOPSYS	- Operating System Name
X.....E	DAY	- Day of Month
X.....E	HOUR	- Hour of Day
X..M..E	MONTH	- Month of Year
X..MY.E	SYSID	- System Identifier

```
 X.....E  WEEK     - Week of Year
 X..MY.E  YEAR     - Year of Century
 X..MY.E  ZONE     - Time Zone

Common Data Elements

 X.....E  AUMAPNUM - Application Number
 X.....E  AUMAPU   - Application Name
 X.....E  AUMGRPID - Group Identification Number
 X.....E  AUMLNTTY - Logon TTY Name
 X.....E  AUMPIN   - Process Identification Number
 X.....E  AUMPPID  - Parent Identification Number
 X.....E  AUMPRI   - Dispatch Priority
 X.....E  AUMPROG  - Program Name
 X.....E  AUMUSER  - User Logon Name
 X.....E  DAYNAME  - Name of Day of Week
 X..MY.E  DURATION - Recording Interval Time
 X..MY.E  ENDTS    - End Time Stamp
 X..MY.E  STARTTS  - Start Time Stamp

Accumulated Data Elements

 X..MY.E  KPSCOST  - Processing Charges
 X..MY.E  KPSCPUTM - Total CPU Use
 X..MY.E  KPSDBKCT - Block I/Os
 X..MY.E  KPSDSKIO - Physical Disk I/Os
 X..MY.E  KPSLDKCT - Logical Disk I/Os
 X..MY.E  KPSPKBCT - Total I/O Transfer Count in KBytes
 X..MY.E  KPSRSSKB - Resident Set Size
 X..MY.E  KPSRUNTM - Total Run Time
 X..MY.E  KPSSCPTM - CPU Wait Time
 X..MY.E  KPSSDKTM - Disk I/O Wait Time
 X..MY.E  KPSSDRTM - Directed to Wait Time
 X..MY.E  KPSSIMTM - Software Impede Wait Time
 X..MY.E  KPSSIPTM - IPC Resource Wait Time
 X..MY.E  KPSSLNTM - LAN I/O Wait Time
 X..MY.E  KPSSNFTM - Network File System Wait Time
 X..MY.E  KPSSOITM - Non-LAN/DISK/TERM I/O Wait Time
 X..MY.E  KPSSSYTM - System Wait Time
 X..MY.E  KPSSVMTM - Virtual Memory Wait Time
 X..MY.E  KPSVSSKB - Virtual Memory Regions

Derived Data Elements

 X..MY.E  KPSAVDRT - Physical Disk I/O Rate
 X..MY.E  KPSAVLRT - Logical Disk I/O Rate
 X..MY.E  KPSDBKRT - Block I/O Rate
 X..MY.E  KPSPKBRT - I/O Transfer Rate in KBytes
```

5.4.3 - AUMKPS Usage Considerations

1. To populate this file, you must ensure that your organization's MeasureWare parm file specifies the SHORTLIVED threshold setting and omits the NOKILLED threshold setting.

2. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

3. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

5.4.4 - AUMKPS Retrieval Examples

The following example lists resources consumed by processes that were ended between 9:00 and 9:59 am. Processes are grouped by its associated application on the report.

```
PROC SORT DATA=DETAIL.AUMKPS01 OUT=KPS;
  BY SYSID AUMAPU ENDTS;
RUN;
PROC PRINT DATA=KPS LABEL NOOBS;
  BY SYSID AUMAPU;
  PAGEBY AUMAPU;
  VAR AUMAPNUM AUMPROG AUMUSER AUMPIN AUMPRI
      KPSRUNTM KPSCPUTM KPSDSKIO KPSLDKCT
  ;
  SUM KPSRUNTM KPSCPUTM KPSDSKIO KPSLDKCT
  ;
  WHERE (HOUR=9);
  TITLE "Detail Killed Processes Resource Consumption";
RUN;
```

5.5 - UNIX Process Exception File (AUMIPS)

The UNIX Process Exception file contains measurements about interesting processes (processes that exceed thresholds) and is useful for performance and problem analyses and modeling. It includes information about CPU use, disk, memory, and transactions as well as reason codes (defining why the process is interesting). The AUMIPS file is derived from MeasureWare Application metrics.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMIPS File Organization
2 - AUMIPS Data Elements List
3 - AUMIPS Usage Considerations
4 - AUMIPS Retrieval Examples

5.5.1 - AUMIPS File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID HOUR AUMPROG	AUMOPSYS ENDTS	YEAR AUMUSER	MONTH AUMLNTTY	DAY AUMPIN
DAYS	N/A				
WEEKS	N/A				
MONTHS	N/A				
YEARS	N/A				
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-5. AUMIPS Timespan Granularity Chart

5.5.2 - AUMIPS Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

- X - DETAIL
- D - DAYS
- W - WEEKS
- M - MONTHS
- Y - YEARS
- T - TABLES AREA
- . - File is not supported

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
Sequence/Summary Data Elements		
X.....E	AUMOPSYS	- Operating System Name
X.....E	AUMPIN	- Process Identification Number
X.....E	DAY	- Day of Month
X.....E	HOUR	- Hour of Day
X.....E	SYSID	- System Identifier
X.....E	WEEK	- Week of Year
X.....E	YEAR	- Year of Century
X.....E	ZONE	- Time Zone

```
+-------------------------------------------------------------------------------------------------------------------------------------+
|                                                                                                                                     |
|                                                                                X.....E  IPSVSSKB - Virtual Memory Regions           |
| Common Data Elements                                                                                                                |
|                                                                                                                                     |
|  X.....E  AUMAPNUM - Application Number                                                                                             |
|  X.....E  AUMAPU   - Application Name                                                                                               |
|  X.....E  AUMGRPID - Group Identification Number                                                                                    |
|  X.....E  AUMLNTTY - Logon TTY Name                                                                                                 |
|  X.....E  AUMPPID  - Parent Identification Number                                                                                   |
|  X.....E  AUMPRI   - Dispatch Priority                                                                                              |
|  X.....E  AUMPROG  - Program Name                                                                                                   |
|  X.....E  AUMUSER  - User Logon Name                                                                                                |
|  X.....E  DAYNAME  - Name of Day of Week                                                                                            |
|  X.....E  DURATION - Recording Interval Time                                                                                        |
|  X.....E  ENDTS    - End Time Stamp                                                                                                 |
|  X.....E  MONTH    - Month of Year                                                                                                  |
|  X.....E  STARTTS  - Start Time Stamp                                                                                               |
|                                                                                                                                     |
| Retained Data Elements                                                                                                              |
|                                                                                                                                     |
|  X.....E  IPSINTR  - Interest Reason Codes                                                                                          |
|  X.....E  IPSPRMID - PRM Group Identification Number                                                                                |
|  X.....E  IPSSTOPR - Blocked State                                                                                                  |
|                                                                                                                                     |
| Accumulated Data Elements                                                                                                           |
|                                                                                                                                     |
|  X.....E  IPSALITM - Process Alive Time                                                                                             |
|  X.....E  IPSCPUTM - Total CPU Use                                                                                                  |
|  X.....E  IPSCSWTM - CPU Use for Context Switching                                                                                  |
|  X.....E  IPSDBKCT - Block I/Os                                                                                                     |
|  X.....E  IPSDBRCT - Block Reads                                                                                                    |
|  X.....E  IPSDBWCT - Block Writes                                                                                                   |
|  X.....E  IPSDFRCT - File System Disk Reads                                                                                         |
|  X.....E  IPSDFWCT - File System Disk Writes                                                                                        |
|  X.....E  IPSDMMCT - Virtual Memory I/Os                                                                                            |
|  X.....E  IPSDSKIO - Physical Disk I/Os                                                                                             |
|  X.....E  IPSDSYCT - File System Management I/Os                                                                                    |
|  X.....E  IPSFSDCT - File System Disk I/Os                                                                                          |
|  X.....E  IPSINTTM - CPU Use for Interrupt Handling                                                                                 |
|  X.....E  IPSLRDCT - Logical Disk Reads                                                                                             |
|  X.....E  IPSLWTCT - Logical Disk Writes                                                                                            |
|  X.....E  IPSNICTM - CPU Use at Nice Priority                                                                                       |
|  X.....E  IPSNORTM - CPU Use at Normal Priority                                                                                     |
|  X.....E  IPSPFMAJ - Page Faults to Disk                                                                                            |
|  X.....E  IPSPFMIN - Page Faults to Memory                                                                                          |
|  X.....E  IPSPKBCT - I/O Transfer Count in KBytes                                                                                   |
|  X.....E  IPSRSSKB - Resident Set Size                                                                                              |
|  X.....E  IPSRTMTM - CPU Use at Real Time Priority                                                                                  |
|  X.....E  IPSRUNTM - Total Accumulated Run Time                                                                                     |
|  X.....E  IPSSSCTM - CPU Use for System Calls                                                                                       |
|  X.....E  IPSSYSTM - CPU Use in System Mode                                                                                         |
|  X.....E  IPSTHCNT - Total Number of Thread Counts                                                                                  |
|  X.....E  IPSUSRTM - CPU Use in User Mode                                                                                           |
|                                                                                                                                     |
|                                                                                                                                     |
|                                                                                                                                     |
+-------------------------------------------------------------------------------------------------------------------------------------+
```

5.5.3 - AUMIPS Usage Considerations

1. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

2. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

5.5.4 - AUMIPS Retrieval Examples

The following sample code lists processes that became interesting between 9:00-9:59 a.m. on an HP-UX system. The report includes information for the interest reason, CPU use, physical I/Os, transactions, and transaction response time.

```
PROC SORT DATA=DETAIL.AUMIPS01 OUT=IPS;
  BY SYSID ENDTS;
RUN;
PROC PRINT DATA=IPS (WHERE=(HOUR EQ 9)) LABEL;
  BY SYSID;
  PAGEBY SYSID;
  VAR ENDTS AUMPROG AUMUSER AUMPIN AUMPRI AUMAPU IPSINTR
      IPSCPUTM IPSDSKIO
  ;
  TITLE "Interesting Processes Summary Report";
RUN;
```

5.6 - UNIX LAN Activity File (AUMNIF)

The UNIX LAN Activity file contains interval-based measurements useful for capacity analysis. It includes information about LAN packets, errors, and collisions, and is derived from MeasureWare NETIF metrics.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMNIF File Organization
2 - AUMNIF Data Elements List
3 - AUMNIF Usage Considerations
4 - AUMNIF Retrieval Examples

5.6.1 - AUMNIF File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	NIFNAME	YEAR	MONTH
	DAY	HOUR	ENDTS		
DAYS	SYSID	AUMOPSYS	NIFNAME	YEAR	MONTH
	DAY	HOUR			
WEEKS	SYSID	AUMOPSYS	NIFNAME	YEAR	WEEK
	ZONE	HOUR			
MONTHS	SYSID	AUMOPSYS	NIFNAME	YEAR	MONTH
	ZONE				
YEARS	SYSID	AUMOPSYS	NIFNAME	YEAR	ZONE
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-6. AUMNIF Timespan Granularity Chart

5.6.2 - AUMNIF Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
Sequence/Summary Data Elements		
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	NIFNAME	- Device Name for LAN Interface
XDWMY.E	SYSID	- System Identifier
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
Common Data Elements		
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
Retained Data Elements		
XDWMY.E	NIFMTUSZ	- MTU Size
XDWMY.E	NIFQUEUE	- Outbound Queue Length
XDWMY.E	NIFSPEED	- Interface Speed
Accumulated Data Elements		
XDWMY.E	NIFCOLCT	- LAN Collisions
XDWMY.E	NIFERRCT	- LAN Errors
XDWMY.E	NIFPKICT	- Inbound LAN Packets
XDWMY.E	NIFPKOCT	- Outbound LAN Packets
Derived Data Elements		
XDWMY.E	NIFPSCOL	- LAN Collision Rate
XDWMY.E	NIFPSERR	- LAN Error Rate
XDWMY.E	NIFPSPKI	- Inbound LAN Packet Rate
XDWMY.E	NIFPSPKO	- Outbound LAN Packet Rate

5.6.3 - AUMNIF Usage Considerations

1. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

2. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

5.6.4 - AUMNIF Retrieval Examples

The following example provides an overview of LAN activity on your systems, ranked by LAN traffic.

```
%LET BY = SYSID YEAR MONTH DAY NIFNAME;
%LET BREAK = NIFNAME;
PROC SORT DATA=DAYS.AUMNIF01 OUT=NIF;
  BY &BY;
RUN;
%MACRO NEGATE;
  PKIORT = SUM(NIFPSPKI,NIFPSPKO);
%MEND;
DATA NIF;
  SET NIF;
  DATE = DATEPART(ENDTS);
  FORMAT DATE MMDDYY8.;
  %NIFFMT(TS=DAYS,DERV=YES);
  %NIFLBL(DERV=YES);
  %NIFSUM;
RUN;
%LET BY = SYSID DATE;
%LET BREAK = DATE;
PROC SORT; BY &BY DESCENDING PKIORT;
DATA NIF2 (KEEP=&BY NIFNAME RANK      NIFPSPKI NIFPSPKO
           NIFPSERR NIFPSCOL)
          ;
  SET NIF;
  BY &BY;
  RETAIN RANK 0;
  IF FIRST.&BREAK THEN RANK = 0;
  RANK+1;
  OUTPUT;
  LABEL RANK = 'Rank';
RUN;
PROC PRINT LABEL;
  BY &BY;
  PAGEBY &BREAK;
  ID RANK;
  VAR  NIFNAME NIFPSPKI NIFPSPKO NIFPSERR NIFPSCOL;
  TITLE "Daily LAN Activity Ranked by Traffic";
RUN;
```

5.7 - UNIX Disk Activity File (AUMDDA)

The UNIX Disk Activity file contains global system measurements useful for capacity analysis. It includes information about logical and physical reads and writes, as well as utilization rates. It is derived from MeasureWare Global metrics.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMDDA File Organization
2 - AUMDDA Data Elements List
3 - AUMDDA Usage Considerations
4 - AUMDDA Retrieval Examples

5.7.1 - AUMDDA File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	DDANAME	YEAR	MONTH
	DAY	HOUR	ENDTS		
DAYS	SYSID	AUMOPSYS	DDANAME	YEAR	MONTH
	DAY	HOUR			
WEEKS	SYSID	AUMOPSYS	DDANAME	YEAR	WEEK
	ZONE	HOUR			
MONTHS	SYSID	AUMOPSYS	DDANAME	YEAR	MONTH
	ZONE				
YEARS	SYSID	AUMOPSYS	DDANAME	YEAR	ZONE
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-7. AUMDDA Timespan Granularity Chart

5.7.2 - AUMDDA Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDWMY.E	DDANAME	- Disk Name
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	SYSID	- System Identifier
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
		Common Data Elements
XDWMY.E	AUMLDEV	- Logical Device Number
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
		Retained Data Elements
XDWMY.E	DDADIR	- File System Directory
		Accumulated Data Elements
XDWMY.E	DDADDPDR	- Time Processing Disk Request
XDWMY.E	DDADFRCT	- File System Reads
XDWMY.E	DDADFWCT	- File System Writes
XDWMY.E	DDADMMCT	- Virtual Memory I/Os
XDWMY.E	DDADRRCT	- Raw Reads
XDWMY.E	DDADRWCT	- Raw Writes
XDWMY.E	DDADSKIO	- Physical I/Os
XDWMY.E	DDADUTTM	- Disk Busy Time
XDWMY.E	DDALRDCT	- Logical Disk Reads
XDWMY.E	DDALWTCT	- Logical Disk Writes
XDWMY.E	DDAPKBCT	- Disk I/O Transfer Count in KBytes
XDWMY.E	DDAPRDCT	- Physical Reads
XDWMY.E	DDAPWTCT	- Physical Writes
XDWMY.E	DDAQLNSM	- Sum of Avg Queue Length
XDWMY.E	DDARKBCT	- KBytes Read
XDWMY.E	DDASYSCT	- File System Management I/Os
XDWMY.E	DDAWKBCT	- KBytes Written
		Derived Data Elements
XDWMY.E	DDAAVQLN	- Avg Queue Length
XDWMY.E	DDADSKRT	- Physical I/O Rate
XDWMY.E	DDAPCDUT	- Percent Disk Utilization
XDWMY.E	DDAPKBRT	- Disk I/O Transfer Rate in KBytes
XDWMY.E	DDAPSFRD	- File System Read Rate
XDWMY.E	DDAPSFWT	- File System Write Rate
XDWMY.E	DDAPSLRD	- Logical Disk Read Rate
XDWMY.E	DDAPSLWT	- Logical Disk Write Rate
XDWMY.E	DDAPSMEM	- Virtual Memory I/O Rate
XDWMY.E	DDAPSRRD	- Raw Read Rate
XDWMY.E	DDAPSRWT	- Raw Write Rate
XDWMY.E	DDAPSSYS	- File System Management I/O Rate
XDWMY.E	DDARKBRT	- Disk Read Transfer Rate in KBytes
XDWMY.E	DDAWKBRT	- Disk Write Transfer Rate in KBytes

5.7.3 - AUMDDA Usage Considerations

1. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

2. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

5.7.4 - AUMDDA Retrieval Examples

List disk utilization, physical and logical I/O rate, and average queue length for disk devices on the system. Disk devices on the report are ranked in descending order of how busy a device was.

```
%LET BY = SYSID YEAR MONTH DAY DDANAME;
%LET BREAK = DDANAME;
PROC SORT DATA=DAYS.AUMDDA01 OUT=DDA;
  BY &BY;
RUN;

%MACRO NEGATE;
  PSRAWIO = SUM(DDAPSRWT,DDAPSRRD);
  PSFILEIO = SUM(DDAPSFWT,DDAPSFRD);
  PSLOGIO = SUM(DDAPSLWT,DDAPSLRD);
  LABEL PSRAWIO = 'Raw I/O Rate'
        PSFILEIO = 'File System I/O Rate'
        PSLOGIO = 'Logical I/O Rate';
%MEND;

DATA DDA;
  SET DDA;
  DATE = DATEPART(ENDTS);
  FORMAT DATE MMDDYY8.;
  %DDAFMT(TS=DAYS,DERV=YES);
  %DDALBL(DERV=YES);
  %DDASUM;
RUN;

%LET BY = SYSID DATE;
%LET BREAK = DATE;
PROC SORT; BY &BY DESCENDING DDAPCDUT;
DATA DDA2 (KEEP=&BY DDANAME DURATION DDAPCDUT
           DDADUTTM DDAAVQLN DDAPSMEM DDAPSSYS PSRAWIO
           PSFILEIO PSLOGIO RANK);
  SET DDA;
  BY &BY;
  RETAIN RANK 0;
  IF FIRST.&BREAK THEN RANK = 0;
  RANK+1;
  OUTPUT;
  LABEL RANK = 'Rank';
  FORMAT DDAPSMEM DDAPSSYS PSRAWIO PSFILEIO PSLOGIO 8.2;
RUN;

PROC PRINT LABEL;
  BY &BY;
  PAGEBY &BREAK;
  ID RANK;
```

```
   VAR  DDANAME DURATION DDAPCDUT DDADUTTM DDAAVQLN
        DDAPSMEM DDAPSSYS PSRAWIO PSFILEIO PSLOGIO
        ;
   TITLE "Daily Disk Activity Ranked by Disk Busy";
 RUN;
```

5.8 - UNIX System and Collector Profile (AUMMCF)

The UNIX System and Collector Profile contains global system measurements useful for problem diagnosis, change tracking, and reporting. It includes information about the UNIX system on which data was collected, and is derived from MeasureWare Configuration metrics.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMMCF File Organization
2 - AUMMCF Data Elements List
3 - AUMMCF Usage Considerations
4 - AUMMCF Retrieval Examples

5.8.1 - AUMMCF File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	YEAR	MONTH	DAY
	HOUR	ENDTS			
DAYS	N/A				
WEEKS	N/A				
MONTHS	SYSID	AUMOPSYS	YEAR	MONTH	ZONE
YEARS	N/A				
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.
NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-8. AUMMCF Timespan Granularity Chart

5.8.2 - AUMMCF Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
Sequence/Summary Data Elements		
X..M..E	AUMOPSYS	- Operating System Name
X.....E	DAY	- Day of Month
X.....E	HOUR	- Hour of Day
X..M..E	MONTH	- Month of Year
X..M..E	SYSID	- System Identifier
X.....E	WEEK	- Week of Year
X..M..E	YEAR	- Year of Century
X..M..E	ZONE	- Time Zone

```
Common Data Elements

 X..M..E  AUMCPMDL - System Model Identification
 X..M..E  AUMLGVER - Measurement Interface Version Number
 X..M..E  AUMOPSLC - Operating System License
 X..M..E  AUMOPSVR - Operating System Release Number
 X..M..E  AUMSCOPE - Scope Collector Name and Version Number
 X..M..E  AUMSYSID - Name of Host System
 X.....E  DAYNAME  - Name of Day of Week
 X..M..E  ENDTS    - End Time Stamp
 X.....E  ORGSYSID - Originating System Identification

Retained Data Elements

 X..M..E  MCFCACHE - Size of Buffer Cache Configured
 X..M..E  MCFCPUMN - Threshold, Percent of CPU
 X..M..E  MCFINODE - Entries in Inode Table on System
 X..M..E  MCFIORMN - Threshold, Physical Disk I/O Rate
 X..M..E  MCFKRNSZ - Kernel Word Size
 X..M..E  MCFLCKAV - File/Record Locks Available
 X..M..E  MCFLGTPS - Data Types Logged by Collector
 X..M..E  MCFLOSHR - Flag - Log Short-Running Processes
 X..M..E  MCFMCMOD - CPU Model
 X..M..E  MCFMEMSW - Swap Space Configured
 X..M..E  MCFMEMSZ - Main Memory Configured
 X..M..E  MCFMEMUS - Main Memory Available for Users
 X..M..E  MCFMSGQU - Message Queues Available
 X..M..E  MCFNCPUS - Number of Processors Configured
 X..M..E  MCFNDSKS - Number of Disk Drives Configured
 X..M..E  MCFNLANS - Number of LAN Interfaces Configured
 X..M..E  MCFNOKIL - Flag - No Terminating Processes
 X..M..E  MCFNONEW - Flag - No Newly Created Processes
 X..M..E  MCFNPROC - Value of NPROC Parameter
 X..M..E  MCFOPSYS - Operating System Name
 X..M..E  MCFPRCMM - Threshold, Virtual Memory
 X..M..E  MCFSEMID - Semaphore Identifiers Available
 X..M..E  MCFSHMSG - Table Memory Segments Available
 X..M..E  MCFTBLUK - Entries Used by Kernel
 X..M..E  MCFWCPMN - Threshold, Percent of CPU Wait
 X..M..E  MCFWDKMN - Threshold, Percent of Disk Wait
 X..M..E  MCFWIMMN - Threshold, Percent of Semaphore Wait
 X..M..E  MCFWMEMN - Threshold, Percent of Memory Wait

Minimum Data Elements

 X..M..E  MCFMNCHC - Minimum Size of Buffer Cache Configured
 X..M..E  MCFMNCPU - Minimum Number of Processors Configured
 X..M..E  MCFMNDSK - Minimum Number of Disk Drives Configured
 X..M..E  MCFMNLAN - Minimum Num of LAN Interfaces Configured
 X..M..E  MCFMNMEM - Minimum Main Memory Configured
 X..M..E  MCFMNMUS - Minimum Main Memory Available for Users
 X..M..E  MCFMNSWS - Minimum Swap Space Configured

Maximum Data Elements

 X..M..E  MCFMXCHC - Maximum Size of Buffer Cache Configured
 X..M..E  MCFMXCPU - Maximum Number of Processors Configured
 X..M..E  MCFMXDSK - Maximum Number of Disk Drives Configured
 X..M..E  MCFMXLAN - Maximum Num of LAN Interfaces Configured
 X..M..E  MCFMXMEM - Maximum Main Memory Configured
 X..M..E  MCFMXMUS - Maximum Main Memory Available for Users
 X..M..E  MCFMXSWS - Maximum Swap Space Configured
```

5.8.3 - AUMMCF Usage Considerations

1. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

2. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

5.8.4 - AUMMCF Retrieval Examples

The following program provides an overview of the profile on your UNIX systems and data collector.

```
PROC PRINT DATA=DETAIL.AUMMCF01 LABEL NOOBS;
  ID SYSID ENDTS;
  VAR AUMSYSID AUMOPSYS AUMOPSVR AUMCPMDL
      MCFNCPUS MCFNDSKS MCFNLANS MCFMEMSZ MCFMEMUS MCFCACHE
      MCFMEMSW MCFLGTPS;
  TITLE 'Overview of UNIX System and Collector Profile';
RUN;
```

5.9 - UNIX Logical Volume Activity File (AUMVLM)

The UNIX Logical Volume Activity file contains interval-based measurements useful for capacity analyses. It includes information about I/O and space for individual logical volumes. It is derived from MeasureWare Volume metrics, which are available from HP-UX systems.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMVLM File Organization
2 - AUMVLM Data Elements List
3 - AUMVLM Usage Considerations
4 - AUMVLM Retrieval Examples

5.9.1 - AUMVLM File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	VLMNAME	YEAR	MONTH
	DAY	HOUR	ENDTS		
DAYS	SYSID	AUMOPSYS	VLMNAME	YEAR	MONTH
	DAY	HOUR			
WEEKS	SYSID	AUMOPSYS	VLMNAME	YEAR	WEEK
	ZONE	HOUR			
MONTHS	SYSID	AUMOPSYS	VLMNAME	YEAR	MONTH
	ZONE				
YEARS	SYSID	AUMOPSYS	VLMNAME	YEAR	ZONE
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-9. AUMVLM Timespan Granularity Chart

5.9.2 - AUMVLM Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	SYSID	- System Identifier
XDWMY.E	VLMNAME	- Logical Volume Group Name
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
		Common Data Elements
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
		Retained Data Elements
XDWMY.E	VLMDIR	- Logical Volume Directory
XDWMY.E	VLMPTHNM	- Mount Point Path Name
		Accumulated Data Elements
XDWMY.E	VLMLRDCT	- Logical Disk Reads
XDWMY.E	VLMLWTCT	- Logical Disk Writes
XDWMY.E	VLMSPUSM	- Sum of Avg Disk Space Used
		Minimum Data Elements
XDWMY.E	VLMMNSPU	- Minimum Pct Disk Space Used
		Maximum Data Elements
XDWMY.E	VLMMXSPU	- Maximum Pct Disk Space Used
		Derived Data Elements
XDWMY.E	VLMPCSPU	- Percent Disk Space Used
XDWMY.E	VLMPSLRD	- Logical Disk Read Rate
XDWMY.E	VLMPSLWT	- Logical Disk Write Rate

5.9.3 - AUMVLM Usage Considerations

1. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

2. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

5.9.4 - AUMVLM Retrieval Examples

List disk space use, logical disk read rate, and logical write rate for logical volumes on the system. Logical volumes on the report are ranked in descending order based on how full a volume was.

```
%LET BY = SYSID YEAR MONTH DAY VLMGROUP VLMPATH;
%LET BREAK = VLMGROUP VLMPATH;
PROC SORT DATA=DAYS.AUMVLM01 OUT=VLM;
  BY &BY;
RUN;

DATA VLM;
  SET VLM;
  DATE = DATEPART(ENDTS);
  FORMAT DATE MMDDYY8.;
  %VLMFMT(TS=DAYS,DERV=YES);
  %VLMLBL(DERV=YES);
  %VLMSUM;
RUN;

%LET BY = SYSID DATE;
%LET BREAK = DATE;
PROC SORT; BY &BY DESCENDING VLMPCSPU;
DATA VLM2 (KEEP=&BY AUMDIR VLMNAME DURATION VLMPCSPU
           VLMPSLRD VLMPSLWT VLMMXSPU VLMMNSPU RANK);
  SET VLM;
  BY &BY;
  RETAIN RANK 0;
  IF FIRST.&BREAK THEN RANK = 0;
  RANK+1;
  OUTPUT;
  LABEL RANK = 'Rank';
RUN;

PROC PRINT LABEL;
  BY &BY;
  PAGEBY &BREAK;
  ID RANK;
  VAR  VLMGROUP VLMPATH AUMDIR DURATION VLMPCSPU VLMMXSPU
       VLMMNSPU VLMPSLRD VLMPSLWT ;
  TITLE1 "Daily Logical Volume Activity";
  TITLE2 "Ranked by Disk Space Used";
RUN;
```

5.10 - UNIX Business Transaction SLO Activity File (AUMTTA)

The UNIX Business Transaction SLO Activity (AUMTTA) file contains measurements useful for managing the service level objectives your business application users require. It includes information about transaction response time and how transactions perform within established SLOs. This file is derived from MeasureWare Transaction metrics.

The following sections describe the file's organization and list the data elements maintained.

1 - AUMTTA File Organization
2 - AUMTTA Data Elements List
3 - AUMTTA Usage Considerations
4 - AUMTTA Retrieval Examples

5.10.1 - AUMTTA File Organization

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity
DETAIL	N/A
DAYS	N/A
WEEKS	N/A
MONTHS	N/A
YEARS	N/A
TABLES	N/A

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

Figure 5-10. AUMTTA Timespan Granularity Chart

5.10.2 - AUMTTA Data Elements List

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

- X - DETAIL
- D - DAYS
- W - WEEKS
- M - MONTHS
- Y - YEARS
- T - TABLES AREA
- . - File is not supported

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
.....E	AUMOPSYS	- Operating System Name
.....E	DAY	- Day of Month
.....E	HOUR	- Hour of Day
.....E	MONTH	- Month of Year
.....E	SYSID	- System Identifier
.....E	TTANAME	- Transaction Name
.....E	WEEK	- Week of Year
.....E	YEAR	- Year of Century
.....E	ZONE	- Time Zone
		Common Data Elements
.....E	DAYNAME	- Name of Day of Week
.....E	DURATION	- Recording Interval Time
.....E	ENDTS	- End Time Stamp
.....E	INTERVLS	- Number of Recording Intervals
.....E	STARTTS	- Start Time Stamp
		Retained Data Elements
.....E	TTAADFMT	- Correlator Address Format
.....E	TTAARMID	- ARM Transaction ID
.....E	TTAARMIF	- ARM Transaction Information
.....E	TTAARMNM	- ARM Application Name
.....E	TTAARMUN	- ARM Transaction User Name
.....E	TTABINCT	- Number of Transaction Bins Defined
.....E	TTACORAD	- Correlator Address
.....E	TTASLOMX	- Transaction Service Level Objective
.....E	TTATFAIL	- Number of Failed Transactions
.....E	TTATRNID	- Transaction Class ID
.....E	TTAUPR01	- Upper Range Response Time for Bin 1
.....E	TTAUPR02	- Upper Range Response Time for Bin 2
.....E	TTAUPR03	- Upper Range Response Time for Bin 3
.....E	TTAUPR04	- Upper Range Response Time for Bin 4
.....E	TTAUPR05	- Upper Range Response Time for Bin 5
.....E	TTAUPR06	- Upper Range Response Time for Bin 6
.....E	TTAUPR07	- Upper Range Response Time for Bin 7
.....E	TTAUPR08	- Upper Range Response Time for Bin 8
.....E	TTAUPR09	- Upper Range Response Time for Bin 9
.....E	TTAUPR10	- Upper Range Response Time for Bin 10
		Accumulated Data Elements
.....E	TTAABORT	- Aborted Transaction Count
.....E	TTAABTTM	- Cumulative Aborted Transaction Time
.....E	TTACPUTM	- Total CPU Consumed by Transaction
.....E	TTADSKIO	- Physical Disk I/Os per Transaction
.....E	TTALGLIO	- Logical I/Os per Transaction
.....E	TTATRANS	- Completed Transaction Count
.....E	TTATRNTM	- Cumulative Transaction Response Time
.....E	TTATRN01	- Completed Transactions in Bin 1
.....E	TTATRN02	- Completed Transactions in Bin 2
.....E	TTATRN03	- Completed Transactions in Bin 3
.....E	TTATRN04	- Completed Transactions in Bin 4
.....E	TTATRN05	- Completed Transactions in Bin 5
.....E	TTATRN06	- Completed Transactions in Bin 6
.....E	TTATRN07	- Completed Transactions in Bin 7
.....E	TTATRN08	- Completed Transactions in Bin 8
.....E	TTATRN09	- Completed Transactions in Bin 9
.....E	TTATRN10	- Completed Transactions in Bin 10

```
 ......E  TTAVSLCT - Violated SLO Transaction Count

Derived Data Elements

 ......E  TTAAVATM - Average Time Per Aborted Trans
 ......E  TTAAVTTM - Average Transaction Response Time
 ......E  TTAPCB01 - Percent Transactions in Bin 1
 ......E  TTAPCB02 - Percent Transactions in Bin 2
 ......E  TTAPCB03 - Percent Transactions in Bin 3
 ......E  TTAPCB04 - Percent Transactions in Bin 4
 ......E  TTAPCB05 - Percent Transactions in Bin 5
 ......E  TTAPCB06 - Percent Transactions in Bin 6
 ......E  TTAPCB07 - Percent Transactions in Bin 7
 ......E  TTAPCB08 - Percent Transactions in Bin 8
 ......E  TTAPCB09 - Percent Transactions in Bin 9
 ......E  TTAPCB10 - Percent Transactions in Bin 10
 ......E  TTAPCVSL - Percent Transactions Violated SLO
```

5.10.3 - AUMTTA Usage Considerations

1. Exercise care when using the special date and time data elements contained in each CA MICS file. As the file's granularity increases in higher timespans, certain fields lose significance and should not be referenced:

   - o HOUR should not be used in MONTHS and YEARS.
   - o DAY and DAYNAME should not be used in WEEKS, MONTHS, or YEARS.
   - o WEEK should not be used in MONTHS or YEARS.
   - o MONTH should not be used in YEARS.

2. The data elements STARTTS and ENDTS have different meanings when used in the DETAIL timespan from when they are used in the DAYS, WEEKS, MONTHS, and YEARS timespans. The ENDTS and STARTTS, when appearing in the higher timespans, indicate the span of time over which the data has been summarized. STARTTS indicates the beginning of the timespan and ENDTS indicates the end of the timespan.

3. To collect MeasureWare Transaction data, you need to prepare your applications and define service level objectives to the HP Transaction Tracker. Please refer to the HP MeasureWare Agent: Transaction Tracker Handbook for details.

4. The AUMTTA file is turned off by default. After you collect and export transaction data, make sure you turn on the AUMTTA file in AUMGENIN and run AUMCGEN to activate the file.

5.10.4 - AUMTTA Retrieval Examples

List the percentage of transactions that meet the service level objective and the distribution of transactions among transaction bins.

```
 %LET BY=SYSID YEAR MONTH DAY HOUR TTANAME;
 %LET BREAK=TTANAME;
 PROC SORT DATA=DAYS.AUMTTA01 OUT=TTA;
  BY &BY;
 RUN;

 DATA TTA;
   SET TTA;
   DATE = DATEPART(ENDTS);
   IF 8 < HOUR < 18;
   %TTAFMT(TS=DAYS,DERV=YES);
   %TTALBL(DERV=YES);
   %TTASUM;
   PCSLO=100-TTAPCVSL;
   LABEL PCSLO='%Transactions Met SLOs';
   FORMAT PCSLO MAPCT.;
 RUN;

 %LET BY=SYSID DATE HOUR;
 %LET BREAK=HOUR;
 PROC PRINT DAT=TTA LABEL;
  BY &BY;
  PAGEBY &BREAK;
  VAR TTANAME PCSLO TTAPCB01 TTAPCBO2 TTAPCB03
      TTAPCBO4 TTAPCB05 TTAPCB06 TTAPCB07
      TTAPCB08 TTAPCBO9 TTAPCB10
      ;
  TITLE "Business Transaction Service Level Activity";
 RUN;
```

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 6

DATA SOURCES

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 6

DATA SOURCES

Chapter 6. DATA SOURCES

The Analyzer for MeasureWare OS Agents supports the following data sources:

- o Hewlett-Packard's MeasureWare for HP-UX Systems, Releases A.09, B.10, and B.11
- o Hewlett-Packard's MeasureWare for HP-UX Systems, IBM RISC System/6000 Systems
- o Sun Microsystems, SunOS 5.4 - 5.6 and Sun Solaris 2.5, 2.5.1, 2.6, 7, 8, and 9

The following sections provide more detail about the Analyzer's source data:

1 - Data Collection Considerations
2 - Record Descriptions

6.1 - Data Collection Considerations

Hewlett-Packard's MeasureWare includes the following data classes:

- o Configuration data class metrics identify the system and collector on which the data was captured, and the response time distributions and process thresholds established, among other measures.
- o Global data class metrics report on system activity.
- o Application data class metrics capture data by user-defined applications, which are pre-summarized business groups of MeasureWare process data.
- o Process data class metrics capture data about each process.
- o Disk data class metrics capture data about each disk.
- o LVolume data class metrics capture data about each logical volume.
- o Transaction data class metrics capture response time about business transactions.
- o NETIF data class metrics capture data about the network interface.

MeasureWare retains data in log files, purging it as the log files fill up. MeasureWare enables you to extract data in export format for transport to another system.

The terms "export" and "extract" have specific meanings in MeasureWare. Extract is a program that helps you manage log files. Export is an extract program command that converts MeasureWare log file data into a format easily accessible outside of MeasureWare.

Data from Hewlett-Packard's MeasureWare populates the following Analyzer files:

This MeasureWare data ...	is written to this raw log...	which populates these Analyzer files.
Configuration	logglob	AUMMCF
Global	logglob	AUMGPR, AUMNIF
Application	logappl	AUMAPP
Process	logproc	AUMKPS, AUMIPS
Disk	logdev	AUMDDA
Volume	logdev	AUMVLM
Transaction	logtran	AUMTTA

Preparing to Collect Data

Your UNIX system's MeasureWare parm file controls the data that MeasureWare collects. It also enables your organization to define applications. The application definition from your UNIX system cannot be changed once application data is created.

If you anticipate using application data in the Analyzer, perhaps to charge for UNIX system resources through the Analyzer's interface to CA MICS Accounting and Chargeback, ensure that your UNIX systems programmer is aware of your needs and reflects them in your organization's MeasureWare parm file.

If you anticipate using killed process data (the AUMKPS file) for accounting and chargeback or workload measurements, ensure that your UNIX systems programmer specifies the following in your organization's MeasureWare parm file:

- o Shortlived processes are logged by coding the threshold setting SHORTLIVED in the parm file.
- o Killed processes are logged, whether or not they exceeded other thresholds, by omitting the threshold setting NOKILLED in the parm file.

If you anticipate using disk data, ensure that your UNIX systems programmer specified the following in your organization's MeasureWare parm file:

- o Disk metrics are kept by coding log device=disk in the parm file for HP MeasureWare.

Sample MeasureWare Parameter (Parm) File

```
log global application process device=disk,lvm
log transaction

size global=10.0, application=10.0, process=20.0
size device=10.0

threshold cpu = 5.0, disk = 5.0, nonew, shortlived

application = xwindows
file = X*,xload,xclock,xterm,hpterm,grmd,softmsg*,vue

application = network
file = nfs*,biod,automount,inetd,snmp*,rpc*,llbd,netfmt
file = portmap
file = rbootd,telnet*,ftp*,*rlogin*,remsh*,rcp,nktl*
file = nvsisr,ttisr
file = lcsp,gcsp,strmen,strweld,vtdaemon

application = memory_management
file = swapper,vhand,syncer,pageout,fsflush

application = other_user_root
user = root
```

A MeasureWare report file (i.e., script) lists the data you want to export. We provide prepared report files for HP/UX, Sun Solaris, and IBM AIX UNIX platforms. Refer to the tables below to determine the appropriate report file to use based on the release of MeasureWare you are executing. These report files are stored in sharedprefix.MICS.HOLD.PARMS:

HP/UX Report Files

MeasureWare Release	Report File
up to C.02.60	AUMRHP10
C.02.60 and above	AUMRHP12

IBM AIX Report Files

MeasureWare Release	Report File
B.02.19 - C.03.40	AUMRAIX
C.03.40 - C.03.84	AUMRAX12
C.04.5 and above	AUMRAX13

Sun Solaris Report Files

MeasureWare Release	Report File
B.02.04.01 - C.03.45	AUMRSOL
C.03.45 and above	AUMRSO12

You should copy the applicable report file to your UNIX system and use it as your MeasureWare report file, rather than creating your own file.

Some Analyzer files may be empty if your site does not export the data types listed in export report files. (If you do not collect some data types and want to avoid empty files, you can deactivate Analyzer files using standard CA MICS file tailoring procedures.)

For example, if your MeasureWare parm file does not specify log=application, then your AUMAPP (UNIX Application Performance) file will be empty. Consult your UNIX systems programmer to ensure that your site is collecting the data needed to populate the Analyzer's files.

An extract command stream, which runs on your UNIX system, retrieves MeasureWare data using the report file definition.

To collect MeasureWare Transaction data, you need to prepare your applications and define service level objectives to the Transaction Tracker. Please refer to the HP MeasureWare Agent: Transaction Tracker Handbook for details.

Example--Extract Command Stream

An extract command stream to export data using a report file looks like this:

```
extract>
configuration on
global detail
application detail
process on                                  < see 1
disk detail
lvolume detail                              < see 2
netif detail
transaction detail                          < see 3
shift all day
report 'reptfile'                           < see 4
output default,purge                        < see 5
export day -1                               < see 6
exit
```

1: If your site is only interested in killed processes, replace 'process on' with 'process killed'. Specifying 'process killed' extracts only processes that terminated in the measurement interval, reducing the size of the database.
2: Logical volume data is only available for HP-UX systems.
3: If Transaction Tracker data is enabled.
4: Replace 'reptfile' with the name of your report file, such as aumrhp09 or aumrhp10.
5: Default writes the output files to xfrdGLOBAL.bin, xfrdAPPLICATION.bin, xfrdPROCESS.bin, xfrdDISK.bin, xfrdLVOLUME.bin, xfrdNETIF.bin, xfrdTRANSACTION.bin, and xfrdCONFIGURATION.bin and purge overwrites files if they already exist. Instead of using the default, we recommend specifying a specific 'filename,' which creates less work when inputting your data sets into DAY089. MeasureWare enables you to name the output files to meet your site's standards. Purge overwrites files if they already exist.
6: Day -1 means that the prior day's data is being exported.

Another way to export data is by invoking the extract program in shell scripts that look like this:

```
extract -v -r aumrhp12 -gapkdczn
        -c transaction detail -f d951215,purge
        -xp d-1
```

where -f d951215 writes the output file to d951215.

IMPORTANT

The -v option is required when using scopeux version B.02.00 under HP-UX 10.0 or 9.0 to circumvent a known MeasureWare software problem.

Example--Extract Command Stream With SUMMARY

You can control the granularity of the data in the report file using the optional SUMMARY= statement. For example, SUMMARY=15 means that data is being collected every 15 minutes. MeasureWare supports SUMMARY= values ranging from 5 to 1440 minutes. For the Analyzer, we recommend using values of 5, 10, 15, 20, or 30. The default summary interval is 5 minutes, however process data is always in 1 minute intervals (independent of the summary interval specified).

An extract command stream that includes the SUMMARY= statement and a report file looks like this:

```
extract>
configuration on
global summary
application summary
process on                                  < see 1
disk summary
lvolume summary                             < see 2
netif summary
transaction summary                         < see 3
shift all day
report 'reptfile'                           < see 4
output default,purge                        < see 5
export day -1                               < see 6
exit
```

1: Process data is always in 1 minute intervals. If your site is only interested in killed processes, replace 'process on' with 'process killed'. Specifying 'process killed' extracts only processes that terminated in the measurement interval, reducing the size of the database.

2: Logical volume data is only available for HP-UX systems.
3: If Transaction Tracker data is enabled.
4: Replace 'reptfile' with the name of your report file, such as aumrhp09 or aumrhp10.
5: Default writes the output files to xfrsGLOBAL.bin, xfrsAPPLICATION.bin, xfrsPROCESS.bin, xfrsDISK.bin, xfrsLVOLUME.bin, and xfrdCONFIGURATION.bin. MeasureWare enables you to name the output files to meet your site's standards. Purge overwrites files if they already exist.
6: Day -1 means that the prior day's data is being exported.

Another way to export data is by invoking the extract program in shell scripts that look like this:

```
extract -v -r aumrhp12 -GApkDcZN
        -c transaction summary -f d951215,purge
        -xp d-1
```

where -f d951215 writes the output file to d951215.

> **IMPORTANT**
>
> The -v option is required when using scopeux version B.02.00 under HP-UX 10.0 or 9.0 to circumvent a known MeasureWare software problem.

Note: If you are using MeasureWare Desktop Agent, application and process data are not supported.

Transferring Data

When the MeasureWare extract program completes, transfer the data to the MVS system on which CA MICS is installed using a product like CA's XCOM file transfer software. Requirements for transfer are:

- o The data must be transferred in BINARY mode.
- o The data must be in variable length, variable blocked, variable block spanned, or undefined length (V, VB, VBS, or U) records after being transferred to MVS.
- o The logical record length of the data on MVS must not exceed 32760.

Extract Command Stream Warnings

- o Configuration data is required for Analyzer processing. The Analyzer will abend if Configuration data is not present in the exported data.
- o Do not use the 'both' option with the global, disk, or application commands in an extract command stream. Using 'both' can result in MeasureWare placing 5-minute detail data with summary data in an exported output file. CA MICS abends when it detects this scenario to prevent database corruption.

6.2 - Record Descriptions

The data records used by the Analyzer are described in the following table.

Data Source	Exported Record Type	Characteristics	Analyzer for MeasureWare OS Agents File
MeasureWare for HP-UX and IBM RISC System/6000 and SunOS/Solaris	Global	System-wide workload measures. Data records are in 5 minute intervals or the interval length you specified for the SUMMARY= statement in your extract command stream.	AUMGPR--UNIX Global Performance AUMNIF--UNIX LAN Activity
	Process	Interesting process measures. Cut every minute.	AUMKPS--UNIX Process Activity AUMIPS--UNIX Process Exception
	Application	User-defined application process measures. Data records are in 5 minute intervals or the interval length you specified for the SUMMARY= statement in your extract command stream.	AUMAPP--UNIX Application Performance
	Configuration	System configuration and parm file information. One record each time the collector program, scopeux, starts, and when the configuration changes. This data is required for the Analyzer's correct operation.	AUMMCF--UNIX System and Configuration Profile
	Disk	Disk utilization measures. Data records are in 5 minute intervals or the interval length you specified for the SUMMARY= statement in your extract command stream.	AUMDDA--UNIX Disk Activity
	Volume	Logical volume measures. Data records are in 5 minute intervals or the interval length you specified for the SUMMARY= statement in your extract command stream.	AUMVLM--UNIX Logical Volume Activity

Figure 6-1. Record Descriptions (Part 1 of 2)

Data Source	Exported Record Type	Characteristics	Analyzer for MeasureWare OS Agents File
	Transaction	Transaction Tracker data. Data records are in 5 minute intervals or the interval length you specified for the SUMMARY= statement in your extract command stream.	AUMTTA--UNIX Business Transaction SLO Activity
	NETIF	Network interface measures. Data records are in 5 minute intervals or the interval length you specified for the SUMMARY= statement in your extract command stream.	AUMNIF--UNIX LAN Activity

Figure 6-1. Record Descriptions (Part 2 of 2)

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 7

PARAMETERS

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 7

PARAMETERS

Chapter 7. PARAMETERS

Defining product parameters requires that you gain an understanding of your data center and its needs, and that you translate that understanding into CA MICS parameters.

This chapter requests that you:

- o Make several policy decisions
- o Complete various worksheets
- o Translate the worksheet entries into the corresponding CA MICS parameter library entries

These activities represent the major portion of the product installation process.

The CA MICS system administrator can use this chapter, along with the PIOM, as a detailed reference when installing or modifying the product. Chapters 2 and 3 of the PIOM document the mechanics of the CA MICS installation process and include checklists that describe each installation step.

If you have a question at any time during your review of the material presented here, contact Technical Support at http://ca.com/support.

This chapter contains information on the following topics:

1 - Environmental Considerations
2 - Complex Level Parameters
3 - Unit Level Parameters
4 - Other Related Parameters

7.1 - Environmental Considerations

Before coding product parameters, you need to know about the UNIX environment(s) at your installation. Before you specify parameters for the product:

- o Review the default options for the parameters to determine their applicability to your site.
- o Review reporting requirements to determine whether or not you need to code MICF inquiries to satisfy your reporting needs.
- o Review database unit specifications to determine which unit or units should include UNIX data.

7.2 - Complex Level Parameters

This section shows you how to specify the complex level parameters that define MW OS Agent Analyzer processing. The following topics are covered:

1 - Analyzer Definition Statements (AUMGENIN)
2 - Account Code Structure (AUMACCT)
3 - Account Code Derivation Routine (AUMACRT)

7.2.1 - Analyzer Definition Statements (AUMGENIN)

A generation definition statement member is provided for the Analyzer for MeasureWare in sharedprefix.MICS.GENLIB(AUMGENIN). Chapter 4 of the CA MICS System Modification Guide describes the contents of GENIN members.

Remember that you need to run AUMCGEN (in sharedprefix.MICS.CNTL) after tailoring AUMGENIN.

7.2.2 - Account Code Structure (AUMACCT)

In CA MICS, data about UNIX activity is stored by account codes in the UNIX Process Activity (AUMKPS) file.

The parameters you specify in sharedprefix.MICS.PARMS(AUMACCT) define the number of account code fields that will be carried in the AUMKPS file, the length of each field, and the SAS long names that are associated with each field. Note that once you have defined these fields, you must provide a routine that assigns them values (see Section 7.2.3).

Preparing to Define Account Codes

Each installation has its own method for associating its UNIX work with the responsible user, projects, or departments. Before defining account codes, it is important that you investigate your installation's accounting standards to:

- o Identify the current organizational coding system (for example, cost center coding system identifying the division, department, project, and employee) and determine how the account codes are specified.

- o Identify if, and how, the codes are verified to ensure that they correspond to valid definitions. You should always validate account codes and assign unidentified or invalid account codes to a special installation overhead account code. This approach provides you with two benefits. First, it groups all invalid codes under the same category and therefore requires less DASD space for storing the Analyzer's files. Second, it enables you to easily determine how much of this overhead activity is taking place.

Consider the following suggestions in determining the number of account codes required to meet your installation's reporting and analysis needs:

- o Account codes are part of the summarization keys for the files that support them. As such, at least one record is generated for each combination of values. Therefore, the higher the number of account codes, the more DASD space the database will require. But at the same time, a higher number of account codes also supports data analysis at a greater detail.

- o If you anticipate needing to expand the account code structure in the future, establish an extra account code now to eliminate the need to retrofit the database later to add the new account code.

Defining Account Codes

Account code fields names are in the form "AUMACTx", where 'x' is the sequential number of the account code field. If three fields are defined, they will be AUMACT1, AUMACT2, and AUMACT3 (in your accounting structure, these might identify division, department, and project, respectively). The sequential number is called the account code field "level" number. There is a maximum of nine levels. A sample account code structure is provided in sharedprefix.MICS.PARMS(AUMACCT).

Figure 7-1 provides a worksheet for collecting the data. Once you have completed this form, code the contents of sharedprefix.MICS.PARMS(AUMACCT) as follows:

- o A separate statement is coded for each account code level.

- o Blank statements are permitted. Comments are coded by beginning the statement with an '*'.

- o Account levels are provided in order, starting with "1".

- o Up to nine levels are permitted with no gaps between the numbers permitted.

- o The statement format is freeform but positional. All parameters are required. The statement format is:

```
level  length  'descriptive title'

   or

level  mask  length 'descriptive title'
```

where:

level = The level of importance of each account code, with level 1 being the most important and the highest level number being the least important. The levels are defined sequentially starting with 1. You can define up to nine levels. This parameter is required.

mask = An optional parameter that can be used to deactivate the account code in specified time-spans. Specifying a timespan mask requires

coding T(.....) as the second parameter in a statement, where each "." represents a file timespan, in the order of DETAIL, DAYS, WEEKS, MONTHS, YEARS, and TABLES. For each timespan, you can specify a Y to indicate that the account code is active or an N to indicate that the account code is inactive.

If this parameter is skipped, it has a default value of T(YYYYYY).

The following rules apply if you code the mask:

- o DETAIL must always be Y. Account codes may not be deactivated in this timespan.
- o If DAYS is N, then WEEKS, MONTHS, and YEARS must also be N. If the account code is inactive in the DAYS timespan, it cannot be active in higher timespans.
- o If MONTHS is N, then YEARS must be N. If the account code is inactive in the MONTHS timespan, it cannot be active in the YEARS timespan.

length = The length of the account code. The length may range from 1 to 30. This parameter is required.

descriptive title = The title that describes the account code. The length of the title is 1 to 40 characters. The title is embedded in quotes (''). This is a required parameter.

A sample AUMACCT member with three levels is illustrated below:

```
1            8 'DIVISION'
2           20 'USER GROUP'
3 T(YYNNNN) 16 'USER NAME'
```

```
+-------------------------------------------------------------------------+
| INSTALLATION PREPARATION WORKSHEET: AUM Account Code Level Definition    |
|                                                                         |
| PARMS Library Member is AUMACCT                                         |
| Reference Sections: 7.2.2                                               |
+-------------------------------------------------------------------------+
|                                                                         |
|  ACCOUNT    FIELD                                                       |
| CODE LEVEL  LENGTH  ACCOUNT CODE LEVEL DESCRIPTION                      |
|   (1-9)                                                                 |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|     _         __    '_________________________________________'        |
|                                                                         |
|                                                                         |
+-------------------------------------------------------------------------+
| ....5...10...15...20...25...30...35...40...45...50...55...60...65...70.. |
+-------------------------------------------------------------------------+
```

Figure 7-1. Account Code Level Definition Worksheet

7.2.3 - Account Code Derivation Routine (AUMACRT)

After you have defined AUMACCT for the number of account codes to be carried in the UNIX Process Activity (AUMKPS) file, your next step is to code the Account Code Derivation Routine (AUMACRT) in sharedprefix.MICS.PARMS. AUMACRT is a SAS routine that derives the account code data elements during the DAY089 step. For example, if you have defined three account codes in AUMACCT, then your AUMACRT exit must contain SAS code to derive the data elements AUMACT1, AUMACT2, and AUMACT3.

You are responsible for testing the accuracy of the exit routine; however, the Analyzer provides a sample exit routine. The worksheet for coding the AUMACRT exit is shown in Figure 7-2.

To limit retention of an AUMACTx element at a particular timespan, your sharedprefix.MICS.PARMS(AUMACCT) member must define that account code with the TIMESPAN mask, T(.....) as described in Section 7.2.2 of this guide.

A sample AUMACRT member is illustrated below:

```
/*********************************************************/
/*                                                       */
/* AUMACRT - ACCOUNT CODE DERIVATION ROUTINE             */
/*                                                       */
/*           This routine is a user coded routine to build */
/*           the account code fields defined by 'AUMACCT'  */
/*           member in sharedprefix.MICS.PARMS. This       */
/*           routine must correctly build each of the      */
/*           defined 'AUMACTX' codes from 1 to n.          */
/*                                                         */
/* SAMPLE CODE:                                            */
/*                                                         */
/*  DIVISION - set to the 1st word in the application name */
/*  USER GROUP - set to the rest of the application name   */
/*  USER NAME  - set to the user id                        */
/*                                                         */
/*  IF APPLICATION NUMBER IS 1, IT IS COUNTED AS OVERHEAD. */
/*                                                         */
/*       IF AUMAPNUM GT 1 THEN DO;                         */
/*         AUMACT1 = SCAN(AUMAPU,1,':-');                  */
/*         I = INDEXC(AUMAPU,':-');                        */
/*         IF I GT 0 THEN                                  */
/*           AUMACT2 = SUBSTR(AUMAPU,I+1);                 */
/*         ELSE                                            */
/*           AUMACT2 = AUMAPU;                             */
/*         AUMACT3 = AUMUSER;                              */
/*       END;                                              */
/*       ELSE DO;                                          */
/*         AUMACT1 = 'OVHD';                               */
/*         AUMACT2 = '****';                               */
/*         AUMACT3 = '****';                               */
/*       END;                                              */
/*                                                         */
/*********************************************************/

/****************************************/
/* SAMPLE ACCOUNT CODE DERIVATION EXIT */
/****************************************/

IF AUMAPNUM GT 1 THEN DO;
  AUMACT1 = '****';
  AUMACT2 = AUMAPU;
END;
ELSE DO;
  AUMACT1 = 'OVHD';
  AUMACT2 = AUMPROG;
END;
AUMACT3 = AUMUSER;
```

```
INSTALLATION PREPARATION WORKSHEET: AUM Account Code Derivation Routine

PARMS Library Member is AUMACRT
Reference Sections: 7.2.3

* VALIDATE FOR VALID ACCOUNT CODES, WHERE POSSIBLE ;

IF { account data is valid } THEN DO;

* BUILD ACCOUNT CODE FIELDS AS IN THE WORKSHEET 7-1

    AUMACT1=field source 1 ;
    ...
    ...
    ...
    AUMACTx=field source n ;
    ____________________________________________________________________
    ____________________________________________________________________
    ____________________________________________________________________
    ____________________________________________________________________
    ____________________________________________________________________

END;
ELSE DO;

* ROUTINE TO BUILD INSTALLATION OVERHEAD ACCOUNT CODES ;

    AUMACT1='overhead category' ;
    ...
    ...
    ...
    AUMACTx='overhead category' ;
    ____________________________________________________________________
    ____________________________________________________________________
    ____________________________________________________________________
    ____________________________________________________________________

END;

....5...10...15...20...25...30...35...40...45...50...55...60...65...70..
```

Figure 7-2. Account Code Derivation Routine Worksheet

7.3 - Unit Level Parameters

This section shows you how to define the unit level parameters for the Analyzer for MeasureWare. These parameter definitions exist for each CA MICS database unit in which the Analyzer is installed. Some of the parameters are used by the Analyzer parameter generation (AUMPGEN) to generate code that will be executed during the DAY089 step. Other parameters are read directly during DAY089 to determine the derivation of CA MICS data elements and handling of error diagnostics.

The following topics are presented:

1 - Processing Options (AUMOPS)
2 - Input DD Statements (INPUTAUM)
3 - Input Needed to Generate System Code (AUMPGEN)
4 - Application Name Derivation Routine (AUMAURT)
5 - Database Space Modeling (DBMODEL)

7.3.1 - Processing Options (AUMOPS)

The AUMOPS member in prefix.MICS.PARMS is used to define the characteristics of the Analyzer for MeasureWare data input into each unit daily update run. In addition, AUMOPS supports various control statements that are specific to the Analyzer and other control statements that control CA MICS common features, such as Internal Step Restart.

Operational statements are stored in the prefix.MICS.PARMS cccOPS member, where ccc is the component identifier, and are incorporated into the CA MICS system by running the prefix.MICS.CNTL(cccPGEN) job.

```
************************************************************
*                                                          *
*  NOTE:  CHANGES to prefix.MICS.PARMS(cccOPS) members     *
*         REQUIRE EXECUTION of prefix.MICS.CNTL(cccPGEN)   *
*         to take effect.                                  *
*                                                          *
*         In addition, any change to parameters that       *
*         impact the DAILY operational job JCL such as,    *
*                                                          *
*         o  changing RESTART NO to RESTART YES,           *
*                                                          *
*         o  WORK parameter changes when RESTART NO is in  *
*            effect,                                       *
*                                                          *
*         o  Specifying TAPEfff (if this product supports  *
*            a DETAIL level TAPE option),                  *
*                                                          *
*         o  or changes to prefix.MICS.PARMS(INPUTccc),    *
*                                                          *
*         will require regeneration of the DAILY job by    *
*         executing prefix.MICS.CNTL(JCLGEND) or by        *
*         specifying DAILY in prefix.MICS.PARMS(JCLGENU)   *
*         and executing prefix.MICS.CNTL(JCLGENU).         *
*                                                          *
*         Refer to the checklist (if provided) for updating *
*         cccOPS parameters and running required generation *
*         jobs.                                            *
*                                                          *
************************************************************
```

The AUMOPS member contains the following statements, described in detail below. A worksheet for recording the statement values for AUMOPS is shown in Section 7.3.1.9.

Required Statements

OPTIONS and PLATFORM

Optional Statements

COMMON
GLOBAL FORCE
VERIFY
WORK
MULTWORK|NOMULT
RESTART
INCRUPDATE

These statements are described in detail below.

7.3.1.1 - OPTIONS Statement

The OPTIONS statement associates a UNIX system name with an ORGSYSID and is required for the MeasureWare data from each UNIX system input to the DAY089 DAILY update step. Each OPTIONS statement must be immediately followed by an PLATFORM statement.

The OPTIONS statement has the following format:

```
OPTIONS orgsysid hostname
```

where:

orgsysid - is the original SYSID used in prefix.MICS.PARMS(SYSID) to identify this UNIX system.

hostname - is the UNIX system name that is used to run this UNIX system. It is specified as a character string that does not exceed 40 characters. If your hostname contains embedded blanks, enclose the hostname in quotation marks. The first 40 characters of the UNIX system name in the configuration metric GBL_SYSTEM_ID is compared to this value during DAY089 processing.

Sample OPTIONS statements:

```
OPTIONS DC01 unixsys1
OPTIONS NYP1 NEW_YORK_PURCHASING_1
OPTIONS DC01 "WASHINGTON D.C."
```

7.3.1.2 - PLATFORM Statement

A PLATFORM statement is required and must follow each OPTIONS statement. It identifies the platform (HPUX, AIX, or Solaris) of the input data and the JCL reference (ddname) used for the data set name containing the raw data.

The PLATFORM has the following format:

```
--------- required -----  --------- optional --------

PLATFORM orgsysid ddname  FORCE VER scope_rel ERRSKIP
```

where:

platform - identifies the vendor platform that generated the MeasureWare data. Platform must be one of the following:

- HPUX - Data from HP-UX systems
- SOLARIS - Data from Sun SPARC systems
- AIX - Data from IBM RS/6000 systems

orgsysid - is the same original SYSID specified on a previous OPTIONS statement.

ddname - is the ddname used to read MeasureWare data from the system. Specify 'ddname' as '*' if you are using a COMMON statement to provide the ddname for this data.

FORCE VER - is the argument that enables you to bypass the scopeux version check built into the Analyzer and process data from unknown MeasureWare Agent releases. By default, the scopeux version check causes daily update step DAY089 to abend if the raw data version is not recognized.

If you use the FORCE VER argument, ensure the metrics from the release you are using are compatible with the metrics in the releases the Analyzer supports. The FORCE VER argument only applies to the particular data file associated with each PLATFORM statement. Refer to Sections 1.4 and 6.1 of this guide for more information about supported releases.

To use the FORCE VER argument, specify

```
FORCE VER scope_rel
```

where

scope_rel is a scopeux release. You must enter this value as it appears in the HP MeasureWare data. If you encounter an abend due to unrecognized scopeux release, MICSLOG message AUM00070 displays the value you would use for scope_rel.

The FORCE VER argument is ignored if a GLOBAL FORCE statement has been coded for the platform associated with this PLATFORM statement. The optional GLOBAL FORCE statement is discussed later in this section.

ERRSKIP - is the argument that enables you to skip the MeasureWare data when input data corruption is detected. DAY089 will continue processing data from ddnames associated with PLATFORM statements. Use this option with caution because data may be missing from the CA MICS database. Always check MICSLOG messages for warnings that some data is not processed due to the ERRSKIP option being enabled. Correct the data problem as soon as possible.

To use the ERRSKIP keyword, specify

```
ERRSKIP
```

Sample PLATFORM statements:

```
HPUX NYP1 AUMMWA02
AIX SYS1 AUMMWA02
SOLARIS SYS1 AUMMWA02
AIX SYS1 AUMMWA02 FORCE VER C.03.58 ERRSKIP
```

7.3.1.3 - COMMON Statement

Use the COMMON statement to identify the JCL reference (ddname) and optional parameters to use when concatenating data sets from multiple UNIX systems. The COMMON statement is followed by OPTIONS and PLATFORM statement pairs.

The COMMON statement has the following format:

```
-------- required ---------  --------- optional ----------

COMMON  datasource  ddname   FORCE VER scope_rel ERRSKIP
```

where

datasource - is HPUX, SOLARIS, or AIX.

ddname - is the ddname used to read MeasureWare data from the system. Specify 'ddname' as '*' if you are using a COMMON statement to provide the ddname for this data.

FORCE VER - is the argument that enables you to bypass the scopeux version check built into the Analyzer and process data from unknown MeasureWare Agent releases. By default, the scopeux version check causes daily update step DAY089 to abend if the raw data version is not recognized.

If you use the FORCE VER argument, ensure the metrics from the release you are using are compatible with the metrics in the releases the Analyzer supports. The FORCE VER argument only applies to the particular data file associated with each PLATFORM statement. Refer to Sections 1.4 and 6.1 of this guide for more information about supported releases.

To use the FORCE VER argument, specify

```
FORCE VER scope_rel
```

where

scope_rel is a scopeux release. You must enter this value as it appears in the HP MeasureWare data. If you encounter an abend due to unrecognized scopeux release, MICSLOG message AUM00070 displays the value you would use for scope_rel.

The FORCE VER argument is ignored if a GLOBAL FORCE statement has been coded for the platform associated with this PLATFORM statement. The optional GLOBAL FORCE statement is discussed later.

ERRSKIP - is the argument that enables you to skip the MeasureWare data when input data corruption is detected. DAY089 will continue processing data from ddnames associated with PLATFORM statements. Use this option with caution because data may be missing from the CA MICS database. Always check MICSLOG messages for warnings that some data is not processed due to the ERRSKIP option being enabled. Correct the data problem as soon as possible.

To use the ERRSKIP keyword, specify

```
ERRSKIP
```

Sample COMMON statement:

```
COMMON AIX  AUMMWA01
COMMON SOLARIS AUMMWA02 FORCE VER SOLC0381
```

7.3.1.4 - GLOBAL FORCE Statement

The GLOBAL FORCE statement allows you to bypass the scopeux version checking built into the Analyzer. By design, the Analyzer abends when raw data is processed from an unrecognized scopeux version. As discussed earlier in this chapter, the COMMON, HPUX, AIX, and SOLARIS statement FORCE VER argument permits you to override the version check for a specific unknown release. The GLOBAL FORCE statement is more generic. It permits you to process data from any unknown scopeux release, based upon platform.

Refer to Sections 1.4 and 6.1 of this guide for more information about supported releases.

The GLOBAL FORCE statement has the following format:

```
    GLOBAL FORCE platform(s)

  where

    platform  - identifies the vendor platform that generated
    --------    the MeasureWare data.  Platform can be any
                combination of:

                HPUX AIX SOLARIS

                or

                ALL

                Specifying ALL is equivalent to coding HPUX
                AIX and SOLARIS.

                When you code a GLOBAL FORCE statement, raw
                data from the platforms specified will bypass
                the standard release validation.  No DAY089
                abend will occur for raw data processed from
                unrecognized scopeux releases, and both an
                informational (AUM00092) and warning
                (AUM00062) message will be output to MICSLOG.
```

Sample GLOBAL FORCE statements:

```
    GLOBAL FORCE AIX
    GLOBAL FORCE HPUX AIX
    GLOBAL FORCE HPUX SOLARIS
    GLOBAL FORCE ALL
```

7.3.1.5 - VERIFY Statement

The VERIFY statement determines what action to take when data from an input ddname, specified on an AUM statement, comes from a system not named on the associated OPTIONS statement for the ORGSYSID during DAY089 processing.

The VERIFY statement has the following format:

```
    VERIFY action

  where

    action    - possible keywords are ABORT, NOABORT, or
    ------      EXCLUDE. If the VERIFY statement is omitted,
                the default is ABORT.

                When ABORT is coded and the system name in
                the data does not match any of the hostnames
                on the OPTIONS statements, a U998 ABEND stops
                DAY089 processing.

                When EXCLUDE is coded and the system name in
                the data does not match any of the hostnames
                on the OPTIONS statements, the data from the
                specific system is discarded.

                If the data is from a ddname specified on a
                COMMON statement, the VERIFY statement has no
                effect.
```

Sample VERIFY statement:

```
    VERIFY ABORT
```

7.3.1.6 - WORK, MULTWORK, and NOMULT Statements

WORK

This statement is optional. It enables sites experiencing either SAS WORK space allocation problems or out of work space conditions during DAYnnn or INCRnnn (where nnn is the job step number), daily or incremental update processing, to allocate multiple WORK files.

You can allocate multiple WORK files for use during the daily and/or incremental update job step. The maximum number of WORK files you can allocate varies by product. These additional work files are used in conjunction with the single work data set allocated by default using the JCLDEF parameters WORKUNIT and WORKSPACE.

Because the individual space allocation requirement for each WORK file is typically much smaller, it is more likely to be satisfied.

To take advantage of multiple WORK files support, edit prefix.MICS.PARMS(cccOPS) and insert a WORK statement as shown below:

```
WORK n data_set_allocation_parameters
```

where n is the number of WORK data sets

```
        NOTE:  The default is zero (0).
               The maximum is nine (9).
```

data_set_allocation_parameters is one or more data set allocation parameters (for example, STORCLAS or SPACE) separated by spaces.

You can also specify the WORK parameter as the following:

```
WORK n XXX pppp ssss
```

where:

```
  n    is the number of WORK data sets
  XXX  is TRK or CYL
  pppp is the primary allocation
  ssss is the secondary allocation
```

Note: When allocating any number of SAS WORK data sets, be aware that one additional SAS WORK data set is automatically allocated to facilitate sorting. For example, if you allocate six SAS WORK data sets, you will actually get seven.

If you omit the data_set_allocation_parameters or the WORK parameter, the work data sets are allocated according to the values you specified for the WORKUNIT and WORKSPACE parameters in prefix.MICS.PARMS(JCLDEF). Use the data_set_allocation_parameters to override this default, either to alter the space allocation or to use System Managed Storage (SMS) parameters to control data set placement and characteristics.

Note: If you allocate insufficient space for the WORK data sets, DAYnnn and/or INCRnnn processing will fail and can only be restarted from the beginning.

Note: If internal step restart is active, you can override the WORK data set allocation parameters at execution-time using the //PARMOVRD facility. For more information about execution-time override of dynamic data set allocation parameters, see the PIOM, section 2.3.6.

Specify data set allocation parameters, separated by blanks, according to SAS LIBNAME statement syntax. If you need multiple lines, repeat the WORK keyword on the continuation line.

WORK accepts the engine/host options documented in the SAS Companion for the z/OS environment, including STORCLAS, UNIT, SPACE, BLKSIZE, DATACLAS, MGMTCLAS, and VOLSER.

Important! Do not specify the DISP parameter.

Example 1:

```
WORK n STORCLAS=MICSTEMP SPACE=(XXX,(pppp,ssss),RLSE)
```

where:

```
  n        - is the number of WORK data sets.
  STORCLAS - specifies a storage class for a new data set.
             The name can have up to 8 characters.
  SPACE    - specifies how much disk space to provide for
             a new data set being allocated.
  XXX      - is TRK or CYL.
  pppp     - is the primary allocation.
  ssss     - is the secondary allocation.
  RLSE     - specifies that free-space should be released
             when the data set is closed.
```

Example 2:

```
WORK n XXX pppp ssss
```

where:

```
  n        - is the number of WORK data sets.
  XXX      - is TRK or CYL.
  pppp     - is the primary allocation.
  ssss     - is the secondary allocation.
```

Example 3 (multiple lines):

```
WORK n STORCLAS=MICSTEMP UNIT=SYSDA
WORK   SPACE=(xxxx,(pppp,ssss),,,ROUND))
```

where:

```
  n        - is the number of WORK data sets.
  STORCLAS - specifies a storage class for a new data set.
             The name can have up to eight characters.
  UNIT     - specifies the generic unit for a new data set.
             The name can have up to eight characters.
  SPACE    - specifies how much disk space to provide for
             a new data set being allocated.
  XXX      - is TRK or CYL.
  pppp     - is the primary allocation.
  ssss     - is the secondary allocation.
```

Note: Since there is some performance impact when using multiple WORK files, you should specify the minimum number of WORK data sets to meet your work space requirements. As a start, try incrementing the number gradually beginning from the default.

WORK Considerations

How Much Space Should You Allocate?

o First Time Implementation of Multiple Work Files

If this is the first time you are implementing multiple work files for this product in this unit, review prefix.MICS.PARMS(JCLDEF) and find the WORKSPACE parameter. It will resemble this sample statement:

```
WORKSPACE      TRK 500 250
```

The value shows the current SAS WORK space allocation for the unit as a single data set. It also serves as the default value used in the unit's DAYnnn daily update (and/or INCRnnn incremental update) step unless you provide a WORK parameter.

To achieve the equivalent work space allocation of WORKSPACE TRK 500 250 using multiple WORK data sets that will collectively share the work space requirements of the daily and/or incremental update step, you could code either one of these:

```
WORK 2 SPACE=(TRK,(250,125))

WORK 5 SPACE=(TRK,(100,50))
```

To determine the total work space, multiply the number of WORK files (n) by the primary (pppp) and secondary (ssss) values specified.

Note: To simplify the example, only the SPACE parameter is shown above. You can follow either with data set allocation parameters like UNIT or STORCLAS as required for your site.

o Adjusting Allocation for Existing Multiple WORK Files

If you have previously implemented multiple WORK file support for this product in this unit, and you want to change either the number of WORK files or the space allocations, examine prefix.MICS.PARMS(cccOPS) and find the existing WORK statement.

- If the existing WORK statement only specifies the number of WORK files but does not contain space allocation information as shown below:

  ```
  WORK 5
  ```

  Then each of the multiple WORK files is allocated using the values from the WORKSPACE parameter of prefix.MICS.PARMS(JCLDEF), as described earlier under First Time Implementation of Multiple Work Files.

  To increase workspace, you can increase the number of WORK files (for example, change WORK 5 to WORK 6,7,8, or 9), increase the space allocation in the WORKSPACE parameter, or do both.

  To decrease workspace, you can decrease the number of WORK files (for example, change WORK 5 to WORK 4,3,2, or 1), decrease the space allocation in the WORKSPACE parameter, or do both.

  You can also elect to explicitly specify the multiple WORK file space allocation by adding the space allocation values directly to the WORK statement. This will remove the link to the prefix.MICS.PARMS(JCLDEF)

WORKSPACE parameter for multiple WORK file space allocation. This is recommended as it serves to clearly document, in one place, how multiple WORK files are allocated.

- If the existing WORK statement does include space allocation as shown in the examples below:

  WORK 5 TRK 200 100

  or

  WORK 5 SPACE=(TRK,(200,100)) STORCLAS=MICSTEMP

  Simply change the values to meet your needs.

  If you need more work space, you can increase the number of WORK files (for example, change WORK 5 to WORK 6,7,8, or 9), increase the space allocation (for example, change TRK 200 100 to TRK 250 120), or do both.

  To decrease work space, you can decrease the number of WORK files (for example, change WORK 5 to WORK 4,3,2, or 1), decrease the space allocation (for example, change TRK 200 100 to TRK 150 80), or do both.

Note: If internal step restart is NOT active (RESTART NO) and you change the WORK parameter, you must:

o Run cccPGEN
o Run JCLGENU for DAILY (to regenerate DAILY) and, if incremental update is enabled, INCRccc

When internal step restart is active, (RESTART YES), then, when you change WORK and run cccPGEN, changes take effect immediately. There is no need to run JCLGENU.

SASWORK

This statement is optional.

The WORK DD statement in the CA MICS procedures allocates a temporary data set where SAS keeps its temporary data files and other items that SAS uses during processing of the current job.

By default, the space allocated is defined in the member prefix.MICS.PARMS(JCLDEF) with the WORKSPACE and WORKUNIT parameters, then generated into all the JCL procedures for a given unit.

With the SASWORK statement you have the option to override this unit-wide definition to specify the space allocation individually for the current step.

The format of the SASWORK statement is:

SASWORK data_set_allocation_parameters

where data_set_allocation_parameters is one or more data set allocation parameters (for example, STORCLAS or SPACE) separated by spaces.

You can also specify the SASWORK parameter as the following:

SASWORK XXX pppp ssss

where:

XXX is TRK or CYL
pppp is the primary allocation
ssss is the secondary allocation

If you omit the data_set_allocation_parameters or the SASWORK statement, the WORK data set is allocated according to the values you specified for the WORKUNIT and WORKSPACE parameters in prefix.MICS.PARMS(JCLDEF). Use the data_set_allocation_parameters to override this default, either to alter the space allocation or to use System Managed Storage (SMS) parameters to control data set placement and characteristics.

Specify data set allocation parameters, separated by blanks, according to SAS LIBNAME statement syntax. If you need multiple lines, repeat the SASWORK keyword on the continuation line.

Example:

SASWORK STORCLAS=MICSTEMP SPACE=(XXX,(pppp,ssss))

where:

STORCLAS - specifies a storage class for a new data set. The name can have up to 8 characters.
SPACE - specifies how much disk space to provide for a new data set being allocated.
XXX - is TRK or CYL.
pppp - is the primary allocation.
ssss - is the secondary allocation.

1 - Change the Number of Work Files

Note: If you change the SASWORK parameter, you must:

o Run cccPGEN
o Run JCLGENU for DAILY (to regenerate DAILY) and, if
   incremental update is enabled, INCRccc

MULTWORK|NOMULT fff fff ... fff

Since multiple work files usage impacts performance, this
product provides these optional parameters so you can
restrict multiple work files usage to only those files having
excessive space requirements.

Note: You can only use one of these optional parameters with
the WORK statement, NOT both.

The MULTWORK parameter restricts the use of multiple WORK
files to ONLY those listed after the MULTWORK keyword.

MULTWORK fff fff ... fff

where fff is the unique three character identifier

If you need multiple lines, repeat the MULTWORK on the
continuation line.

The NOMULT parameter forces the use of multiple WORK files
for all files EXCEPT those specified after the NOMULT
keyword.

NOMULT fff fff ... fff

where fff is the unique three character identifier

If you need multiple lines, repeat the NOMULT on the
continuation line.

The default is MULTWORK IPS KPS if neither MULTWORK nor
NOMULT parameters are specified.

The following files are eligible for multiple WORK support:

AUM MeasureWare Files

  IPS   UNIX Process Exception File
  KPS   UNIX Process Activity File

The following section discusses changing the WORK option:

7.3.1.6.1 - Change the Number of Work Files

To change the number of work files used in CA MICS Analyzer for MeasureWare processing in Step DAY089, follow the checklist provided below for each unit.

```
*****************************************************
*                                                   *
*           CHANGE NUMBER OF WORK FILES             *
*                                                   *
*****************************************************
```

___ 1. Update the WORK statement in prefix.MICS.PARMS(cccOPS), where (ccc) is the component identifier, to specify the number of work data sets required. Below is an example:

```
WORK n STORCLAS=MICSTEMP SPACE=(XXX,(pppp,ssss))
```

where:

```
n        - is the number of WORK data sets.
STORCLAS - specifies a storage class for a new data
           set.  The name can have up to eight
           characters.
SPACE    - specifies how much disk space to provide
           for a new data set being allocated.
XXX      - is TRK or CYL.
pppp     - is the primary allocation.
ssss     - is the secondary allocation.
```

You should specify the minimum number of WORK data sets to meet your work space requirements. As a start, try incrementing the number gradually beginning from the default.

___ 2. If this is the first time you are implementing multiple work files for this product, then continue with Step 2. If you are just changing the number currently in use, or simply the space definitions, then proceed to Step 3 of this checklist.

Browse sharedprefix.MICS.PROTOLIB(DYcccnnn) and sharedprefix.MICS.PROTOLIB(cccINCR), where (nnn) is the job step number and (ccc) is the product ID for this product, checking for the presence of the WORK symbolic on the EXEC statement to determine if you have previously modified this product to increase the allocation of SAS WORK space.

___ 2a. If you find a WORK symbolic, simply divide the primary and secondary allocation values from the WORK symbolic by the number of work files specified above (value of n on the WORK statement coded in Step 1).

Coding the resulting values will yield the same aggregate space allocation you have been using with a single WORK file. To double your available WORK space, carry out the division, double the results and use the values in the WORK definition above.

___ 2b. If you did not find a WORK symbolic in PROTOLIB, examine prefix.MICS.PARMS(JCLDEF) for each CA MICS unit containing this product. Find the WORKSPACE keyword. The space allocation specified is used for a single SAS WORK file. Perform the same division as described in the previous paragraph to determine the quantity that will yield equivalent total allocation with multiple WORK files. Then adjust the values upward to meet your needs.

___ 3. Submit the job in prefix.MICS.CNTL(cccPGEN).

___ 4. If you specified RESTART YES in the product's cccOPS, you are done. Otherwise, you must do Steps 5, 6, and 7.

___ 5. Edit prefix.MICS.PARMS(JCLGENU) so that it contains a single line that reads:

DAILY

or, if incremental update is enabled for this product in this unit database, specify:

DAILY INCRccc

where ccc is the product ID.

___ 6. Submit the job in prefix.MICS.CNTL(JCLGENU). Ensure that there are no error messages in MICSLOG or SYSTSPRT, that the MICSLOG contains the normal termination message, BAS10999I, and that the job completes with a condition code of zero.

___ 7. The following operational job(s) have changed:

DAILY

INCRccc (if incremental update is enabled)

If your site has implemented the operational CA MICS processes in a scheduling product, the JCL may have to be refreshed in that product. See the scheduling

product's administrator for the exact processes involved in updating that product's representation of the CA MICS jobs.

7.3.1.7 - Internal Step Restart Statements

RESTART YES/NO

This statement is optional. Specify this to activate internal step restart for this product's DAILY and/or INCRccc database update job steps:

RESTART YES

If you do not specify or enable the RESTART parameter, then this option defaults to the following and internal step restart is disabled:

RESTART NO

```
*************************************************************
*                                                           *
*  Note:  Changing the RESTART parameter (either from NO    *
*         to YES or from YES to NO) requires regeneration   *
*         of the DAILY operational job by executing         *
*         prefix.MICS.CNTL(JCLGEND) or by specifying        *
*         DAILY in prefix.MICS.PARMS(JCLGENU) and           *
*         executing prefix.MICS.CNTL(JCLGENU).              *
*                                                           *
*         If incremental update is active for this product, *
*         you must also regenerate the INCRccc job.         *
*                                                           *
*************************************************************
```

Internal step restart can significantly reduce time and resource usage to recover from daily and/or incremental update processing failures. CA MICS uses a checkpoint/restart technique.

- When internal step restart is activated, the database update job step "checkpoints" (or saves) intermediate results (work file contents) and the operational environment at the end of each processing phase.

- Then, if required, the database update step can resume execution at the beginning of the processing phase in which the failure occurred.

- Restart is accomplished by restoring the operational environment from the last checkpoint, bypassing completed processing phases, and resuming execution using intermediate results (work files) from the last checkpoint.

Note: When you activate internal step restart (RESTART YES), the following optional restart parameters are enabled. These parameters have no effect if restart is disabled (RESTART NO). For more details, see the individual parameter descriptions later in this section.

- o RESTARTCKPT data_set_allocation_parameters
- o RESTARTWORK data_set_allocation_parameters
- o DYNAMWAIT minutes

Processing Phases:

This product employs two database update processing phases followed by the two common roll-up phases.

Phase	Description
FORMAT	Read raw input data, convert to SAS format, and output to intermediate work files.
DBUPDATE	Sort intermediate work file contents, eliminate duplicate input data, prepare for DETAIL cycle creation, merge data across optional multiple work files, enhance data content, and create the new DETAIL cycle.
DYSUM	Summarize DETAIL data to create new DAYS cycles and to update current week-to-date and month-to-date cycles.
DYAGE	Cutover new database cycles to production and "age" existing cycles.

RESTART Considerations

- o Overhead

  Enabling internal step restart adds some overhead to the database update job step -- the cost of taking checkpoints and managing saved materials. Since this overhead is relatively constant and independent of input data volume, you may find that costs outweigh potential savings when input data volume is low, for example in a test unit. For high volume, production units, internal step restart support overhead should be a minor portion of total resource usage.

- o Cataloged Work Files

  When internal step restart is enabled, the SAS work data set, internal step restart control data set, and multiple work file data sets are allocated and cataloged with permanent dataset names so they will be retained for use in restart if the step abends. These data sets are deleted when the step completes successfully.

  Prior to enabling internal step restart support, these data sets were probably allocated on system "scratch" space with a temporary, system assigned data set names. If your installation standards do not allow "permanent" data sets on DASD volumes used for temporary work space, you may need to use the WORK, RESTARTCKPT, and RESTARTWORK parameters to direct the internal step restart data sets to a generic unit or storage class that allows cataloged data sets.

- o Dynamic Allocation

  When internal step restart is active, dynamic allocation is employed for the work data sets. If your installation restricts dynamic allocation of large, cataloged data sets, you may need to use the WORK, RESTARTCKPT, and RESTARTWORK parameters to direct work data set allocation to a generic unit or storage class where dynamic allocation is allowed.

- o Data Set Names

  The SAS work data set, internal step restart control data set, and multiple work file data sets are allocated and cataloged according to the standard CA MICS unit database data set name conventions. The default DDNAME and data set names are:

```
      o SAS work data set,
            //cccXWORK DD DSN=prefix.MICS.cccXWORK,.....

      o Internal step restart control data set,
            //cccXCKPT DD DSN=prefix.MICS.cccXCKPT,.....

      o Multiple work file data sets,
            //WORKnn   DD DSN=prefix.MICS.cccWRKnn,.....

    Since these data sets conform to the same data set name
    conventions as your existing CA MICS data sets, there
    should be few, if any, data set name related allocation
    issues.  However, it is possible to override the data set
    names if required.  Please contact CA MICS Product
    Support for assistance if you must alter data set names.


RESTARTCKPT
-----------

This statement is optional.  Specify the following to
override default data set allocation parameters for the
internal step restart checkpoint data set:

RESTARTCKPT   data_set_allocation_parameters

Note:  RESTARTCKPT is ignored when you specify RESTART NO.

The internal step restart checkpoint data set (or cccXCKPT
data set) contains processing status, control, and SAS
environmental information for internal step restart
processing checkpoints.  This includes a copy of the SAS WORK
format and macro catalogs, current macro variable values, and
a description of work files that may be needed to restart
DAYnnn processing.

By default, the cccXCKPT data set is allocated according to
the values you specified for the WORKUNIT and WORKSPACE
parameters in prefix.MICS.PARMS(JCLDEF).  Specify RESTARTCKPT
to override this default, either to alter the space
allocation or to use System Managed Storage (SMS) parameters
to control data set placement and characteristics.

Note:  If you allocate insufficient space for the cccXCKPT
data set, DAYnnn processing will fail and can only be
restarted from the beginning.

Note:  You can override the RESTARTCKPT data set allocation
parameters at execution-time using the //PARMOVRD facility.
For more information about execution-time override of dynamic
data set allocation parameters, see the PIOM, section 2.3.6.

Specify data set allocation parameters, separated by blanks,
according to SAS LIBNAME statement syntax.  If you need
multiple lines, repeat the RESTARTCKPT keyword on the
continuation line.

RESTARTCKPT accepts the engine/host options documented in the
SAS Companion for the z/OS Environment, including STORCLAS,
UNIT, SPACE, BLKSIZE, DATACLAS, MGMTCLAS, and VOLSER.

Important!  DO NOT SPECIFY THE DISP PARAMETER.

Example 1:

RESTARTCKPT  STORCLAS=MICSTEMP SPACE=(xxxx,(pp,ss),,,ROUND)

where:

  STORCLAS - specifies a storage class for a new data set.
             The name can have up to eight characters.

  SPACE    - specifies how much disk space to provide for
             a new data set being allocated, where:

             xxxx is TRK, CYL, or blklen
             pp is the primary allocation
             ss is the secondary allocation

             and ROUND specifies that the allocated space be
             "rounded" to a cylinder boundary when the unit
             specified was a block length.  ROUND is ignored
             with the TRK or CYL options.

Example 2 (multiple lines):

RESTARTCKPT  STORCLAS=MICSTEMP UNIT=SYSDA
RESTARTCKPT  SPACE=(xxxx,(pp,ss),,,ROUND)

where:

  STORCLAS - specifies a storage class for a new data set.
             The name can have up to eight characters.

  UNIT     - specifies the generic unit for a new data set.
             The name can have up to eight characters.

  SPACE    - specifies how much disk space to provide for
             a new data set being allocated.


RESTARTWORK
-----------
```

This statement is optional. Specify the following to override default data set allocation parameters for the internal step restart WORK data set:

```
RESTARTWORK   data_set_allocation_parameters
```

Note: RESTARTWORK is ignored when you specify RESTART NO.

The internal step restart WORK data set (or cccXWORK data set) contains the intermediate work files that are not enabled to multiple work file support, including those files you may have specified on the optional NOMULT statement.

By default, the cccXWORK data set is allocated according to the values you specified for the WORKUNIT and WORKSPACE parameters in prefix.MICS.PARMS(JCLDEF). Specify RESTARTWORK to override this default, either to alter the space allocation or to use System Managed Storage (SMS) parameters to control data set placement and characteristics.

Note: If you allocate insufficient space for the cccXWORK data set, DAYnnn processing will fail and can only be restarted from the beginning.

Note: You can override the RESTARTWORK data set allocation parameters at execution-time using the //PARMOVRD facility. For more information about execution-time override of dynamic data set allocation parameters, see the PIOM, section 2.3.6.

Specify data set allocation parameters, separated by blanks, according to SAS LIBNAME statement syntax. If you need multiple lines, repeat the RESTARTWORK keyword on the continuation line.

RESTARTWORK accepts the engine/host options documented in "SAS Companion for the z/OS Environment", including STORCLAS, UNIT, SPACE, BLKSIZE, DATACLAS, MGMTCLAS, and VOLSER.

Important! DO NOT SPECIFY THE DISP PARAMETER.

Example 1:

```
  RESTARTWORK  STORCLAS=MICSTEMP SPACE=(xxxx,(pp,ss),,,ROUND)
```

where:

```
  STORCLAS - specifies a storage class for a new data set.
             The name can have up to eight characters.

  SPACE    - specifies how much disk space to provide for
             a new data set being allocated, where:

             xxxx is TRK, CYL, or blklen
             pp is the primary allocation
             ss is the secondary allocation

             and ROUND specifies that the allocated space be
             "rounded" to a cylinder boundary when the unit
             specified was a block length.  ROUND is ignored
             with the TRK or CYL options.
```

Example 2 (multiple lines):

```
  RESTARTWORK  STORCLAS=MICSTEMP UNIT=SYSDA
  RESTARTWORK  SPACE=(xxxx,(pp,ss),,,ROUND)
```

where:

```
  STORCLAS - specifies a storage class for a new data set.
             The name can have up to eight characters.

  UNIT     - specifies the generic unit for a new data set.
             The name can have up to 8 characters.

  SPACE    - specifies how much disk space to provide for
             a new data set being allocated.
```

The following section discusses enabling this option:

1 - Enable Internal Step Restart

7.3.1.7.1 - Enable Internal Step Restart

To enable the internal step restart in the CA MICS Analyzer for MeasureWare, follow the checklist provided below for each unit:

```
*****************************************************
*                                                   *
*           ENABLE INTERNAL STEP RESTART            *
*                                                   *
*****************************************************
```

___ 1. Edit prefix.MICS.PARMS(cccOPS), where (ccc) is the component identifier, and specify:

RESTART YES

For additional information on related topic, review the documentation for this product on WORK, RESTARTWORK, and RESTARTCKPT parameters to override default data set allocation parameters.

___ 2. Submit the job in prefix.MICS.CNTL(cccPGEN).

___ 3. Edit prefix.MICS.PARMS(JCLGENU) so that it contains a single line that reads:

DAILY

or, if incremental update is enabled for this product in this unit database, specify:

DAILY INCRccc

where ccc is the product ID.

___ 4. Submit the job in prefix.MICS.CNTL(JCLGENU). Ensure that there are no error messages in MICSLOG or SYSTSPRT, that the MICSLOG contains the normal termination message, BAS10999I, and that the job completes with a condition code of zero.

___ 5. The following operational job(s) have changed:

DAILY

INCRccc (if incremental update is enabled)

If your site has implemented the operational CA MICS processes in a scheduling product, the JCL may have to be refreshed in that product. See the scheduling product's administrator for the exact processes involved in updating that product's representation of the CA MICS jobs.

7.3.1.8 - Incremental Update Statements

INCRUPDATE

This statement is optional. Specify this to enable incremental update for this product:

```
INCRUPDATE YES
```

If you do not specify or enable the INCRUPDATE parameter, then this option defaults to this and incremental update is disabled:

```
INCRUPDATE NO
```

```
*************************************************************
*                                                           *
*  Note:  Changing the INCRUPDATE parameter (either from NO *
*         to YES or from YES to NO) requires regeneration   *
*         of the DAILY operational job by executing         *
*         prefix.MICS.CNTL(JCLGEND) or by specifying        *
*         DAILY in prefix.MICS.PARMS(JCLGENU) and           *
*         executing prefix.MICS.CNTL(JCLGENU).              *
*                                                           *
*         If you specify INCRUPDATE YES, you must also      *
*         generate the INCRccc, cccIUALC, and cccIUGDG jobs *
*         (where ccc is the 3 character product ID).        *
*         Depending on the options you select, you may also *
*         need to execute the cccIUALC and/or cccIUGDG      *
*         jobs.                                             *
*                                                           *
*************************************************************
```

Incremental update can significantly reduce time and resource usage in the DAILY job by letting you split out a major portion of daily database update processing into multiple, smaller, incremental updates executed throughout the day.

- o Standard CA MICS database update processing involves (1) reading and processing raw input data to generate DETAIL and DAYS level CA MICS database files, followed by (2) summarization of DETAIL/DAYS level data to update week-to-date and month-to-date database files.

- o When you activate incremental update:

  - You can execute the first-stage processing (raw data input to create DETAIL/DAYS files) multiple times throughout the day, each time processing a subset of the total day's input data.

  - Then, during the final update of the day (in the DAILY job), the incremental DETAIL/DAYS files are "rolled-up" to the database DETAIL and DAYS timespans, and then summarized to update the week-to-date and month-to-date files.

- o Incremental update is independent of your internal step restart or DBSPLIT specifications. You have the option to perform incremental updates with or without internal step restart support.

- o Incremental update is activated and operates independently by product. The incremental update job for this product, INCRccc (where ccc is the product ID), can execute concurrently with the incremental update job for another product in the same unit database.

- o The CA MICS database remains available for reporting and analysis during INCRccc job execution.

```
*************************************************************
*                                                           *
*  Note:  CA MICS is a highly configurable system           *
*         supporting up to 36 unit databases, each of which *
*         can be configured and updated independently.      *
*         Incremental update is just one of the options you *
*         can use to configure your CA MICS complex.        *
*                                                           *
*         All efforts should be made to employ CA MICS      *
*         configuration capabilities to minimize issues     *
*         prior to activating incremental update.  For      *
*         example:                                          *
*                                                           *
*         o  Splitting work to multiple units is an         *
*            effective way to enable parallel database      *
*            update processing                              *
*                                                           *
*         o  Adjusting account code definitions to ensure   *
*            adequate data granularity while minimizing     *
*            total database space and processing time       *
*                                                           *
*         o  Tailoring the database to drop measurements    *
*            and metrics of lesser value to your            *
*            data center, thereby reducing database update  *
*            processing and resource consumption            *
*                                                           *
*         While incremental update is intended to reduce    *
*         DAILY job elapsed time, total resource usage of   *
*         the combined INCRccc and DAILY jobs steps can     *
*         increase due to the additional processing         *
*         required to maintain the incremental update       *
*         "to-date" files and for roll-up to the unit       *
*         database.  The increased total resource usage     *
*         will be more noticeable with small data volumes,  *
*         where processing code compile time is a greater   *
*         percentage of total processing cost.              *
*                                                           *
*************************************************************
```

Note: When you activate incremental update (INCRUPDATE YES), the following optional incremental update parameters are enabled. These parameters have no effect if incremental update is disabled (INCRUPDATE NO). For more details, see the individual parameter descriptions later in this section.

o INCRDB PERM/TAPE/DYNAM

o INCRDETAIL data_set_allocation_parameters

o INCRDAYS data_set_allocation_parameters

o INCRCKPT data_set_allocation_parameters

o INCRSPLIT USE/IGNORE data_set_allocation_parameters

Incremental update processing reads and processes raw measurement data to create and maintain DETAIL and DAYS level "to-date" files for the current day.

o These incremental update database files are maintained on unique z/OS data sets, independent of the standard CA MICS database files, and independent of any other product's incremental update database files. There is one data set each for DETAIL and DAYS level "to-date" data and a single incremental update checkpoint data set for this product in this unit.

o The incremental update DETAIL and DAYS files can be permanent DASD data sets, or they can be allocated dynamically as needed and deleted after DAILY job processing completes. Optionally, you can keep the incremental update DETAIL and DAYS files on tape, with the data being loaded onto temporary DASD space as needed for incremental update or DAILY job processing. See the INCRDB PERM/TAPE/DYNAM option for more information.

After activating incremental update, you will use three incremental update facility jobs found in prefix.MICS.CNTL (Note that ccc is the product ID):

o cccIUALC

   You execute this job to allocate and initialize the incremental update checkpoint file, and optionally the incremental update DETAIL and DAYS database files. cccIUALC is generally executed just ONE time.

o cccIUGDG

   You execute this job to add generation data group (GDG) index definitions to your system catalog in support of the INCRDB TAPE option. cccIUGDG is generally executed just ONE time.

o INCRccc

   This is the job you execute for each incremental update. You will integrate this job into your database update procedures for execution one or more times per day to process portions of the total day's measurement data.

   Note: The DAILY job is run once at the end of the day. It will perform the final incremental update for the day's

data, and then roll-up the incremental DETAIL/DAYS files to the database DETAIL and DAYS timespans and update the week-to-date and month-to-date files.

INCRUPDATE Considerations

o Overhead

Incremental update is intended to reduce DAILY job resource consumption and elapsed time by offloading a major portion of database update processing to one or more executions of the INCRccc job. In meeting this objective, incremental update adds processing in the INCRccc and DAILY jobs to accumulate data from each incremental update execution into the composite "to-date" DETAIL and DAYS incremental update files, and also adds processing in the DAILY job to copy the incremental update files to the unit database DETAIL and DAYS timespans. The amount of this overhead and the savings in the DAILY job are site-dependent, and will vary based on input data volume and on the number of times INCRccc is executed each day.

In addition, activating incremental update will cause additional compile-based CPU time to be consumed in the DAYnnn DAILY job step. The increase in compile time is due to additional code included for each file structure in support of the feature. This increase should be static based on the scope of the CA MICS data integration product in terms of files. This compile-time increase does not imply an increase in elapsed or execution time. Incremental update allows I/O bound, intensive processing (raw data inputting, initial CA MICS transformation, etc.) to be distributed outside of the DAILY job. I/O processing is the largest contributor to elapsed time in large volume applications. Thus, the expected overall impact is a decrease in the actual runtime of the DAYnnn job step.

o Increased "Prime Time" Workload

By offloading work from the DAILY job to one or more INCRccc executions throughout the day, you are potentially moving system workload and DASD work space usage from the "off-hours," (when the DAILY job is normally executed) to periods of the day where your system resources are in highest demand. You should schedule INCRccc executions carefully to avoid adverse impact to batch or online workloads. For example, if your site's "prime shift" is 8:00 AM to 5:00 PM, you might choose to schedule incremental updates for 7:00 AM (just before "prime shift") and 6:00 PM (just after "prime shift"), with the DAILY job executing just after midnight.

o Increased DASD Usage

The DASD space required for the incremental update DETAIL and DAYS database files is in addition to the DASD space already reserved for the CA MICS database. By default, the incremental update database files are permanently allocated, making this DASD space unavailable for other applications. In general, you can assume that the incremental update database files will require space equivalent to two cycles of this product's DETAIL and DAYS timespan files.

Alternatively, the incremental update database files can be allocated in the first incremental update of the day and deleted by the DAILY job (see the INCRDB DYNAM option later in this section). This approach reduces the amount of time that the DASD space is dedicated to incremental update, and lets the amount of DASD space consumed increase through the day as you execute each incremental update.

A third option is to store the incremental update database files on tape (see the INCRDB TAPE option). With this approach, the DASD space is required just for the time that each incremental update or DAILY job step is executing. Note that while this alternative reduces the "permanent" DASD space requirement, the total amount of DASD space required while the incremental update or DAILY jobs are executing is unchanged. In addition, the TAPE option adds processing to copy the incremental update files to tape, and to reload the files from tape to disk.

Note: The incremental update checkpoint file is always a permanently allocated disk data set. This is a small data set and should not be an issue.

o Operational Complexity

Incremental update expands your measurement data management and job scheduling issues. You must ensure that each incremental update and the DAILY job processes your measurement data chronologically; that is, each job must see data that is newer than the data processed by the prior job. By incrementally updating the database, you have more opportunities to miss a log file, or to process a log out of order.

o Interval End Effects

Each incremental update processes a subset of the day's measurement data, taking advantage of early availability of some of the day's data, for example, when a measurement log fills and switches to a new volume. This can cause a problem if the measurement log split occurs while the data source is logging records for the end of a measurement interval, thus splitting the data for a single measurement interval across two log files. When an incremental update processes the first log file, the checkpoint high end timestamp is set to indicate that this split measurement interval has been processed. Then, when the rest of the measurement interval's data is encountered in a later update, it can be dropped as duplicate data (because data for this measurement interval end timestamp has already been processed).

Appropriate scheduling of log dumps and incremental updates can avoid this problem. For example, if you plan to run incremental updates at 7:00 AM and 6:00 PM, you could force a log dump in the middle of the measurement interval just prior to the scheduled incremental update executions. This is an extension of the procedure you may already be using for end-of-day measurement log processing. The objective is to ensure that all records for each monitor interval are processed in the same incremental update.

o Dynamic Allocation

When you activate incremental update and specify TAPE or DYNAM for the INCRDB parameter, dynamic allocation is employed for the incremental update database files. If your site restricts dynamic allocation of large, cataloged data sets, you must use the INCRDETAIL and INCRDAYS parameters to direct incremental update data set allocation to a generic unit or storage class where dynamic allocation is allowed.

o Data Set Names

The incremental update database files are allocated and cataloged according to standard CA MICS unit database data set name conventions. The DDNAME and default data set names are (where ccc is the product ID):

o Incremental update checkpoint file,

```
//IUCKPT   DD DSN=prefix.MICS.ccc.IUCKPT,.....
```

o Incremental update DETAIL

```
//IUDETAIL DD DSN=prefix.MICS.ccc.IUDETAIL,.....
```

o Incremental update DAYS

```
//IUDAYS   DD DSN=prefix.MICS.ccc.IUDAYS,....
```

Since these data sets conform to the same data set name conventions as your existing CA MICS data sets, there should be few, if any, data-set-name-related allocation issues. However, it is possible to override the data set names if required. Contact Technical Support at http://ca.com/support for assistance if you must change data set names.

INCRDB

This statement is optional. The default is this:

```
INCRDB PERM
```

Note: INCRDB is ignored when you specify INCRUPDATE NO.

Specify this statement or take the default, to keep the incremental update database DETAIL and DAYS files on permanently allocated DASD data sets:

```
INCRDB PERM
```

Execute the prefix.MICS.CNTL(cccIUALC) job to allocate the incremental update database files.

```
************************************************************
*                                                          *
*  Note:  The incremental update checkpoint file is always *
*         a permanently allocated DASD data set.           *
*                                                          *
************************************************************
```

Specify this to offload the incremental update DETAIL and DAYS files to tape between incremental update executions:

```
INCRDB TAPE #gdgs UNIT=name
```

With the TAPE option, the incremental update DETAIL and DAYS DASD data sets are dynamically allocated at the beginning of the incremental update job or DAILY job step, and then are deleted after the job step completes.

- o The first incremental update job of the day allocates and initializes the incremental update database files. At the end of the job, the DETAIL and DAYS files are copied to a new (+1) generation of the incremental update tape data sets. Then the DASD files are deleted.

- o Subsequent incremental update jobs restore the DASD incremental update database files from the current, (0) generation, incremental update tape data sets before processing the input measurement data. At the end of the job, the DETAIL and DAYS files are copied to a new (+1) generation of the incremental update tape data sets. Then the DASD files are deleted.

- o The DAILY job step also restores the DASD incremental update database files from the (0) generation tape files before processing the input data, but does NOT copy the incremental update database files to tape. Thus, the DAILY job actually creates a new, null (+1) generation.

- o Use the #gdgs parameter to specify the maximum number of incremental update tape generations. The minimum is 2 and the maximum is 99, with a default of 5. Set the number of generations equal to or greater than the number of incremental updates, including the DAILY job you plan to execute each day. This facilitates restart and recovery if you encounter problems requiring you to reprocess portions of the daily measurement data.

- o Use the optional UNIT=name parameter to specify a tape unit name for the incremental update database output tapes. The default is to use the same tape unit as the input tapes.

- o A special index must be created in your system catalog for each of the incremental update tape data set generation data groups. The prefix.MICS.CNTL(cccIUGDG) job will generate the statements to create the incremental update GDG index definitions.

  - Before each index is built, it is deleted. These DLTX (or DELETE) statements causes an error message if no entry exists. This is done so that you can change the number of entries without having to delete each of the index entries.

  - DLTX and BLDG (or DELETE and DEFINE) fail if there is a cataloged data set with the same index. IDCAMS issues a message and gives a return code of 8. This issue is not a problem for non-GDG entries or if the GDG already has the desired number of entries.

  - If you want to change the number of entries kept in a GDG with cataloged data sets, do the following:

    1. Uncatalog any existing entries in the GDG.
    2. Delete the index with a DLTX (or DELETE).
    3. Create the index with a BLDG (or DEFINE).
    4. Catalog any entries that are uncataloged in step 1.

- o The incremental update tape data set names are as follows, where ccc is the product ID:

  - Incremental update tape DETAIL file
    tapeprefix.MICS.ccc.IUXTAPE.GnnnnV00

  - Incremental update tape DAYS file
    tapeprefix.MICS.ccc.IUDTAPE.GnnnnV00

```
*************************************************************
*                                                           *
*  Note:  The INCRDETAIL and INCRDAYS parameters are        *
*         required when you specify INCRDB TAPE.            *
*                                                           *
*************************************************************
```

Specify this parameter to allocate dynamically the incremental update DETAIL and DAYS DASD data sets in the first incremental update of the day, and then delete these data sets at the end of the DAILY job step:

INCRDB DYNAM

- o With this option, no space is used for the incremental update database files during the time between the end of the DAILY job step and the beginning of the next day's first incremental update.
- o With this approach, you can set the data set allocation parameters so that the incremental update DETAIL and DAYS data sets start out with a minimum allocation and then grow through secondary allocations as more space is required for subsequent incremental updates. For example, enough space for one incremental update.

```
*************************************************************
*                                                           *
*  Note:  The INCRDETAIL and INCRDAYS parameters are        *
*         required when you specify INCRDB DYNAM.           *
*                                                           *
*************************************************************
```

INCRDETAIL

This statement is required if you specify either of these:

INCRDB TAPE

INCRDB DYNAM

Otherwise, this statement is optional. There is no default.

Specify this to define data set allocation parameters for the incremental update DETAIL data set (IUDETAIL):

INCRDETAIL data_set_allocation_parameters

Note: INCRDETAIL is ignored when you specify INCRUPDATE NO.

The incremental update DETAIL data set (IUDETAIL) contains the current incremental update detail-level database files, and the DETAIL "to-date" data for the current daily update cycle. You should allocate DASD space equivalent to two cycles of this product's DETAIL timespan data.

If you specified INCRDB PERM (the default), your INCRDETAIL parameter specifications are used in generating the cccIUALC job (where ccc is the product ID).

- o You will execute the cccIUALC job to allocate and initialize the incremental update database and checkpoint files.
- o Omit the INCRDETAIL parameter if you prefer to specify data set allocation parameters directly in the generated prefix.MICS.CNTL(cccIUALC) job.

If you specified INCRDB TAPE or INCRDB DYNAM, your INCRDETAIL parameter specifications are used in incremental update DETAIL data set dynamic allocation during incremental update or DAILY job step execution.

- o The INCRDETAIL parameter is required for the TAPE or DYNAM option.
- o Specify data set allocation parameters, separated by blanks, according to SAS LIBNAME statement syntax. If you need multiple lines, repeat the INCRDETAIL keyword on the continuation line.
- o INCRDETAIL accepts the engine/host options documented in the SAS Companion for the z/OS Environment, including STORCLAS, UNIT, SPACE, BLKSIZE, DATACLAS, MGMTCLAS, and VOLSER.

  Important! DO NOT SPECIFY THE DISP PARAMETER.
- o You can override the INCRDETAIL data set allocation parameters at execution-time using the //PARMOVRD facility. For more information about execution-time override of dynamic data set allocation parameters, see the PIOM, section 2.3.6.

Example 1:

INCRDETAIL STORCLAS=MICSTEMP SPACE=(xxxx,(pp,ss),,,ROUND)

where:

STORCLAS - specifies a storage class for a new data set. The name can have up to eight characters.

SPACE - specifies how much disk space to provide for a new data set being allocated, where:

xxxx is TRK, CYL, or blklen
pp is the primary allocation
ss is the secondary allocation

and ROUND specifies that the allocated space be "rounded" to a cylinder boundary when the unit specified was a block length. ROUND is ignored with the TRK or CYL options.

Example 2 (multiple lines):

```
INCRDETAIL   STORCLAS=MICSTEMP UNIT=SYSDA
INCRDETAIL   SPACE=(xxxx,(pp,ss),,,ROUND)
```

where:

STORCLAS - specifies a storage class for a new data set. The name can have up to eight characters.

UNIT - specifies the generic unit for a new data set. The name can have up to eight characters.

SPACE - specifies how much disk space to provide for a new data set being allocated.

INCRDAYS

This statement is required if you specify either of these:

INCRDB TAPE

INCRDB DYNAM

Otherwise, this statement is optional. There is no default.

Specify this to define data set allocation parameters for the incremental update DAYS data set (IUDAYS):

```
INCRDAYS      data_set_allocation_parameters
```

Note: INCRDAYS is ignored when you specify INCRUPDATE NO.

The incremental update DAYS data set (IUDAYS) contains the current incremental update days-level database files, and the DAYS "to-date" data for the current daily update cycle. You should allocate DASD space equivalent to two cycles of this product's DAYS timespan data.

If you specified INCRDB PERM (the default), your INCRDAYS parameter specifications are used in generating the cccIUALC job (where ccc is the product ID).

- o You will execute the cccIUALC job to allocate and initialize the incremental update database and checkpoint files.
- o Omit the INCRDAYS parameter if you prefer to specify data set allocation parameters directly in the generated prefix.MICS.CNTL(cccIUALC) job.

If you specified INCRDB TAPE or INCRDB DYNAM, your INCRDAYS parameter specifications are used in incremental update DAYS data set dynamic allocation during incremental update or DAILY job step execution.

- o The INCRDAYS parameter is required for the TAPE or DYNAM option.
- o Specify data set allocation parameters, separated by blanks, according to SAS LIBNAME statement syntax. If you need multiple lines, repeat the INCRDAYS keyword on the continuation line.
- o INCRDAYS accepts the engine/host options documented in the SAS Companion for the z/OS Environment, including STORCLAS, UNIT, SPACE, BLKSIZE, DATACLAS, MGMTCLAS, and VOLSER.

  Important! DO NOT SPECIFY THE DISP PARAMETER.
- o You can override the INCRDAYS data set allocation parameters at execution-time using the //PARMOVRD facility. For more information about execution-time override of dynamic data set allocation parameters, see the PIOM, Section 2.3.6.

Example 1:

```
INCRDAYS     STORCLAS=MICSTEMP SPACE=(xxxx,(pp,ss),,,ROUND)
```

where:

STORCLAS - specifies a storage class for a new data set.

```
              The name can have up to eight characters.

  SPACE    - specifies how much disk space to provide for
             a new data set being allocated, where:

             xxxx is TRK, CYL, or blklen
             pp   is the primary allocation
             ss   is the secondary allocation

             and ROUND specifies that the allocated space be
             "rounded" to a cylinder boundary when the unit
             specified was a block length.  ROUND is ignored
             with the TRK or CYL options.

Example 2 (multiple lines):

INCRDAYS     STORCLAS=MICSTEMP UNIT=SYSDA
INCRDAYS     SPACE=(xxxx,(pp,ss),,,ROUND)

where:

  STORCLAS - specifies a storage class for a new data set.
             The name can have up to eight characters.

  UNIT     - specifies the generic unit for a new data set.
             The name can have up to eight characters.

  SPACE    - specifies how much disk space to provide for
             a new data set being allocated.


INCRCKPT
--------

This statement is optional.  Specify this to override default
data set allocation parameters for the incremental update
checkpoint data set:

INCRCKPT      data_set_allocation_parameters

Note:  INCRCKPT is ignored when you specify INCRUPDATE NO.


The incremental update checkpoint data set tracks incremental
update job status and the data that has been processed during
the current daily update cycle.  The incremental update
checkpoint is used to detect and block the input of duplicate
data during incremental update processing.  This data set
will be exactly the same size as prefix.MICS.CHECKPT.DATA
(the unit checkpoint data set), usually 20K to 200K depending
on the prefix.MICS.PARMS(SITE) CKPTCNT parameter (100-9999).
```

Your INCRCKPT parameter specifications are used in generating the cccIUALC job (where ccc is the product ID).

- o You will execute the cccIUALC job to allocate and initialize the incremental update checkpoint file. If you specified INCRDB PERM, then the cccIUALC job will also allocate the incremental update DETAIL and DAYS database files.
- o By default the incremental update checkpoint data set is allocated as SPACE=(TRK,(5,2)) using the value you specified for the prefix.MICS.PARMS(JCLDEF) DASDUNIT parameter.
- o Omit the INCRCKPT parameter if you prefer to override data set allocation parameters directly in the generated prefix.MICS.CNTL(cccIUALC) job.

Specify data set allocation parameters, separated by blanks, according to SAS LIBNAME statement syntax. If you need multiple lines, repeat the INCRCKPT keyword on the continuation line.

INCRCKPT accepts the engine/host options documented in the SAS Companion for the MVS Environment, including STORCLAS, UNIT, SPACE, BLKSIZE, DATACLAS, MGMTCLAS, and VOLSER.

Important! DO NOT SPECIFY THE DISP PARAMETER.

Example 1:

```
INCRCKPT     STORCLAS=MICSTEMP SPACE=(xxxx,(pp,ss),,,ROUND)
```

where:

```
  STORCLAS - specifies a storage class for a new data set.
             The name can have up to eight characters.

  SPACE    - specifies how much disk space to provide for
             a new data set being allocated, where:

             xxxx is TRK, CYL, or blklen
             pp   is the primary allocation
             ss   is the secondary allocation

             and ROUND specifies that the allocated space be
             "rounded" to a cylinder boundary when the unit
             specified was a block length.  ROUND is ignored
             with the TRK or CYL options.
```

Example 2 (multiple lines):

```
INCRCKPT     STORCLAS=MICSTEMP UNIT=SYSDA
INCRCKPT     SPACE=(xxxx,(pp,ss),,,ROUND)
```

where:

STORCLAS - specifies a storage class for a new data set. The name can have up to eight characters.

UNIT - specifies the generic unit for a new data set. The name can have up to eight characters.

SPACE - specifies how much disk space to provide for a new data set being allocated.

INCRSPLIT

This statement is optional and defaults to this:

```
INCRSPLIT IGNORE
```

Specify the following if you want the incremental update job for this product to get input measurement data from the output of the SPLITSMF job. The optional data_set_allocation_parameters are used by the SPLITSMF job when creating the measurement data file for this product.

```
INCRSPLIT  USE  data_set_allocation_parameters
```

Note: INCRSPLIT is ignored when you specify INCRUPDATE NO.

This option would be used when multiple products in a single unit database are enabled to incremental update. The SPLITSMF job performs the same function for incremental update jobs as the DAILY job DAYSMF step performs for the DAYnnn database update steps.

- o The SPLITSMF job dynamically allocates, catalogs, and populates prefix.MICS.ccc.IUSPLTDS data sets for each product in the unit database for which you specified both the INCRUPDATE YES and INCRSPLIT USE parameters. These data sets are then deleted after processing by the appropriate INCRccc job.
- o Specify data set allocation parameters, separated by blanks, according to SAS LIBNAME statement syntax. If you need multiple lines, repeat the INCRSPLIT keyword on each continuation line.
- o INCRSPLIT accepts the engine/host options documented in the SAS Companion for the MVS Environment, including STORCLAS, UNIT, SPACE, BLKSIZE, DATACLAS, MGMTCLAS, and VOLSER.

  Important! DO NOT SPECIFY THE DISP PARAMETER.

Specify the following or accept the default if you want the incremental update jobs for this product to get their input measurement data from the data sets specified in the INPUTccc (or INPUTSMF) member of prefix.MICS.PARMS:

```
INCRSPLIT  IGNORE
```

When you specify INCRSPLIT IGNORE, this product will NOT participate in SPLITSMF job processing.

Example 1:

```
INCRSPLIT USE  STORCLAS=MICSTEMP SPACE=(xxxx,(pp,ss),,,ROUND)
```

where:

STORCLAS - specifies a storage class for a new data set. The name can have up to eight characters.

SPACE - specifies how much disk space to provide for a new data set being allocated, where:

xxxx is TRK, CYL, or blklen
pp is the primary allocation
ss is the secondary allocation

and ROUND specifies that the allocated space be "rounded" to a cylinder boundary when the unit specified was a block length. ROUND is ignored with the TRK or CYL options.

Example 2 (multiple lines):

```
INCRSPLIT  USE  STORCLAS=MICSTEMP UNIT=SYSDA
INCRSPLIT       SPACE=(xxxx,(pp,ss),,,ROUND)
```

where:

```
STORCLAS - specifies a storage class for a new data set.
           The name can have up to eight characters.

UNIT     - specifies the generic unit for a new data set.
           The name can have up to eight characters.

SPACE    - specifies how much disk space to provide for
           a new data set being allocated.
```

DYNAMWAIT

This statement is optional. Specify the following:

```
     DYNAMWAIT   minutes
```

to override the default amount of time, in minutes, the DAILY and/or INCRccc job will wait for an unavailable data set.

Note: This optional parameter is not normally specified. The system default is adequate for most data centers.

Internal Step Restart and Incremental Update facilities use z/OS dynamic allocation services to create new data sets and to access existing data sets. Data set naming conventions and internal program structure are designed to minimize data set contention. However, if data set allocation does fail because another batch job or online user is already using a data set, DAILY and/or INCRccc processing will wait 15 seconds and then try the allocation again. By default, the allocation will be attempted every 15 seconds for up to 15 minutes. After 15 minutes, the DAILY or INCRccc job will abort.

If data set contention in your data center does cause frequent DAILY or INCRccc job failures, and you are unable to resolve the contention through scheduling changes, you may want to use the DYNAMWAIT parameter to increase the maximum number of minutes the DAILY and/or INCRccc jobs will wait for the data set to become available.

On the other hand, if your data center standards require that the DAILY and/or INCRccc jobs fail immediately if required data sets are unavailable, specify the following:

```
     DYNAMWAIT 0
```

Note: You can override the DYNAMWAIT parameter at execution-time using the //PARMOVRD facility. For more information about execution-time override of dynamic data set allocation parameters, see the PIOM, section 2.3.6.

The following section discusses enabling this option:

1 - Implement Incremental Update

7.3.1.8.1 - Implement Incremental Update

To implement incremental update in the CA MICS Analyzer for MeasureWare, follow the checklist provided below:

```
*****************************************************
*                                                   *
*           IMPLEMENT INCREMENTAL UPDATE            *
*                                                   *
*****************************************************
```

___ 1. Edit prefix.MICS.PARMS(cccOPS), where (ccc) is the component identifier:

- o Specify the following:

  INCRUPDATE YES

- o If you want to store the incremental update database files on tape between incremental updates, specify this:

  INCRDB TAPE #gdgs

- o If you want to allocate the incremental update database files during the first incremental update of the day and delete these data sets at the end of the DAILY job step, specify this:

  INCRDB DYNAM

- o If you specified INCRDB TAPE or INCRDB DYNAM, then you must also specify this:

  INCRDETAIL data_set_allocation_parameters
  INCRDAYS data_set_allocation_parameters

- o If you want the incremental update job for this product to get input measurement data from the output of the SPLITSMF job, specify this:

  INCRSPLIT USE data_set_allocation_parameters

- o For additional information on related topic, review the documentation for this product on INCRCKPT, INCRDETAIL, INCRDAYS, or INCRSPLIT parameters to override default data set allocation parameters.

___ 2. Submit the job in prefix.MICS.CNTL(cccPGEN).

___ 3. Edit prefix.MICS.PARMS(JCLGENU) so that it contains two or more lines reading:

DAILY
INCRccc cccIUALC cccIUGDG

___ 4. Submit the job in prefix.MICS.CNTL(JCLGENU). Ensure that there are no error messages in MICSLOG or SYSTSPRT, that the MICSLOG contains the normal termination message, BAS10999I, and that the job completes with a condition code of zero.

___ 5. Edit the job in prefix.MICS.CNTL(cccIUALC).

- o Inspect and/or specify data set allocation parameters for the incremental update database and checkpoint files. If you specified INCRDB TAPE or INCRDB DYNAM, the cccIUALC job will only allocate the incremental update checkpoint data set.

- o Submit the job. Ensure that there are no error messages in MICSLOG or SASLOG, and that the job completes with a condition code of zero.

___ 6. If you specified INCRDB TAPE, submit the job in prefix.MICS.CNTL(cccIUGDG) to define generation group indexes for the incremental update DETAIL and DAYS tape data sets. Examine SASLOG, MICSLOG, and SYSPRINT to verify that the generation group indexes were correctly defined.

Note: You may see error messages for the DLTX (or DELETE) statements. This is not a problem. cccIUGDG deletes each index prior to defining it, and an error message is issued if the index does not yet exist (e.g., if this is the first time you ran the cccIUGDG job).

___ 7. The following operational job(s) have changed:

DAILY INCRccc

If your site has implemented the operational CA MICS processes in a scheduling product, the JCL may have to be refreshed in that product. See the scheduling product's administrator for the exact processes involved in updating that product's representation of the CA MICS jobs.

___ 8. Implement operational procedures for gathering input measurement data and executing incremental updates (INCRccc) during the day.

You may also need to modify operational procedures for the DAILY job to ensure that processing is limited to input measurement data that has not been input to one of the day's incremental update executions.

7.3.1.9 - Sample AUMOPS and Examples

Following is a sample of required and optional statements that may be coded in prefix.MICS.PARMS(SNTOPS).

```
* NO COMMON STATEMENT IN USE
* DEFINE ONE OPTIONS STATEMENT FOR EACH UNIX SYSTEM

 VERIFY ABORT
 GLOBAL FORCE AIX

 OPTIONS       CHID CHICAGO_DEMO
   HPUX        CHID AUMMWA01

 OPTIONS       CHIP CHICAGO_PRODUCTION
   HPUX        CHIP AUMMWA02

 OPTIONS       TRP1 KIPLING
   AIX         TRP1 AUMAIX03


* NO COMMON STATEMENT IN USE
* USING THE FORCE VER KEYWORD

 OPTIONS       CHIP CHICAGO_PRODUCTION
   HPUX        CHIP AUMMWA02  FORCE VER C.03.7


* COMMON STATEMENT IN USE
* DEFINE ONE OPTIONS STATEMENT FOR EACH UNIX SYSTEM

 VERIFY ABORT
 COMMON HPUX   AUMMWA01

 OPTIONS       CHID CHICAGO_DEMO
   HPUX        CHID AUMMWA02

                                          -----+
 OPTIONS       NYP1 NEW_YORK_PURCHASING_1  |
   HPUX        NYP1 *                      | THESE OPTIONS
                                           | ARE ASSOCIATED
 OPTIONS       NYP2 NEW_YORK_PURCHASING_2  | WITH THE COMMON
   HPUX        NYP2 *                      | STATEMENT
                                          -----+
 OPTIONS       CHIP CHICAGO_PRODUCTION
   HPUX        CHIP AUMMWA03
```

```
* COMMON STATEMENT IN USE
* USING BOTH FORCE VER & ERRSKIP KEYWORDS

 COMMON HPUX   AUMMWA01  FORCE VER C.03.7  ERRSKIP

                                      -----+
 OPTIONS       NYP1 NEW_YORK_PURCHASING_1  |
   HPUX        NYP1 *                      | THESE OPTIONS
                                           | ARE ASSOCIATED
 OPTIONS       NYP2 NEW_YORK_PURCHASING_2  | WITH THE COMMON
   HPUX        NYP2 *                      | STATEMENT
                                      -----+
 OPTIONS       CHIP CHICAGO_PRODUCTION
   HPUX        CHIP AUMMWA03
```

Figure 7-3 provides a worksheet for collecting the data.

```
                              +-------------------------------------------------------------------------+
                              | INSTALLATION PREPARATION WORKSHEET: AUM Unit Options Definitions         |
                              |                                                                         |
                              | PARMS Library Member is AUMOPS                                          |
                              +-------------------------------------------------------------------------+
                              |                                                                         |
                              | AUM PROCESSING OPTIONS:                                                 |
                              |                                                                         |
                              |                                                                         |
  General-unit options        |    VERIFY       _______(ABORT/NOABORT/EXCLUDE)                          |
                              |                                                                         |
                              |    GLOBAL FORCE _______(HPUX/AIX/SOLARIS/ALL)                           |
                              |                                                                         |
                              |    WORK n       _______(data_set_allocation_parameters)                 |
                              |                                                                         |
                              |    RESTART      YES/NO                                                  |
                              |                                                                         |
                              |    INCRUPDATE   YES/NO                                                  |
                              |    INCRDB       PERM/TAPE/DYNAM                                         |
                              |    INCRDETAIL   data_set_allocation_parameters                          |
                              |    ...                                                                  |
                              |                                                                         |
                              |                                                                         |
  Special-purpose option      |    COMMON       ________ ________ ___________________ _______           |
                              |               (data source) (ddname)   (FORCE option)    (ERRSKIP)      |
                              |                                                                         |
                              |                                                                         |
  System-oriented options     |    OPTIONS      _________ ________                                      |
                              |                (orgsysid) (system name)                                 |
                              |                                                                         |
                              |    HPUX         ________ ________ ___________________ _______           |
                              |                (orgsysid)  (ddname)   (FORCE option)    (ERRSKIP)       |
                              |                                                                         |
                              |    AIX          ________ ________ ___________________ _______           |
                              |                (orgsysid)  (ddname)   (FORCE option)    (ERRSKIP)       |
                              |                                                                         |
                              |    SOLARIS      ________ ________ ___________________ _______           |
                              |                (orgsysid)  (ddname)   (FORCE option)    (ERRSKIP)       |
                              |                                                                         |
                              +-------------------------------------------------------------------------+
                              | ....5...10...15...20...25...30...35...40...45...50...55...60...65...70...75|
                              +-------------------------------------------------------------------------+
```

Figure 7-3. AUMOPS Options Definition Worksheet

7.3.2 - Input DD Statements (INPUTAUM)

The Analyzer for MeasureWare processes input files from MeasureWare platform data. You must move the data from your UNIX system to the MVS system that runs CA MICS.

You supply the input DD statements that define the UNIX input files in prefix.MICS.PARMS(INPUTAUM). The DD statements are referenced by the MWA/HPUX and COMMON statements of AUMOPS.

There must be one DD statement in the INPUTAUM member for each ddname identified in prefix.MICS.PARMS(AUMOPS).

Remember that ddnames are defined in the AUMOPS member in prefix.MICS.PARMS and that the DCB characteristics for transferred data are required only for nonlabelled tapes. DCB characteristics are included in "Transferring Data" in Section 6.1 of this guide.

The INPUTAUM format is

```
//ddname   DD DISP=SHR,DSN=your.dataset.name
```

where

ddname is a ddname specified on the MWA/HPUX or COMMON statements in prefix.MICS.PARMS(AUMOPS).

your.dataset.name is the name of the data set that contains the input data.

Sample INPUTAUM statements:

```
//AUMMWA01 DD DISP=SHR,DSN=PREFIX.MWA-HPUX.CHICAGO1
//AUMMWA02 DD DISP=SHR,DSN=PREFIX.MWA-HPUX.CHICAGO2
```

Providing Data to CA MICS

The CA MICS system administrator is responsible for getting the data to be processed each day into the MVS data sets you specify here. The preferred method is using some form of file transmission such as our XCOM software or some similar facility that allows you to transfer files between your UNIX systems and your MVS systems.

A worksheet for coding INPUTAUM is shown in Figure 7-4.

```
+-------------------------------------------------------------------------+
| INSTALLATION PREPARATION WORKSHEET: INPUTAUM JCL Definitions            |
|                                                                         |
| PARMS Library Member is INPUTAUM                                        |
| Reference Section:  7.3.2, CA MICS Analyzer for MeasureWare Guide       |
+-------------------------------------------------------------------------+
|                                                                         |
|                                                                         |
| This definition is required to specify the DD statements for UNIX       |
| data that will be read by the CA MICS DAILY job.                        |
|                                                                         |
|      //@ SAMPLE INPUTAUM                                                |
|      //@                                                                |
|      //@                                                                |
|      //________ DD DISP=SHR,DSN=_____________________________________   |
|      //________ DD DISP=SHR,DSN=_____________________________________   |
|      //________ DD DISP=SHR,DSN=_____________________________________   |
|      //@                                                                |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
|                                                                         |
+-------------------------------------------------------------------------+


NOTE: Because of the text substitution process used by JCLGEN, to produce
      '//*', '/*', and '&' in the JCL produced by JCLGEN, the following
      symbols must be used instead:

       //@    produces    //*
       /@     produces    /*
       ?      produces    &
```

Figure 7-4. INPUTAUM JCL Definitions Worksheet

7.3.3 - Input Needed to Generate System Code (AUMPGEN)

Some CA MICS code for this CA MICS Analyzer is generated at the unit level through the AUMPGEN job. The prefix.MICS.PARMS members that must be completed prior to AUMPGEN execution are:

- o SYSID (Computing system parameters); remember to execute BASPGEN as described in Section 7.4 of this guide.
- o AUMOPS (Processing options definition)
- o INPUTAUM (Input DD statements)

Figure 7-5 identifies the prefix.MICS.PARMS input members, output members, and SAS macro names used for the unit-level code generation process. A member may contain the definition of more than one SAS macro. The output libraries are at the database unit level (prefix).

prefix.MICS.PARMS Input Member	prefix.MICS.MUOLIB Input Format	USER.SOURCE Output Members and SAS Macro Names
AUMOPS, INPUTAUM	$SYSID $AUMSYS	$AUMMSTR %AUMMWAIN

Figure 7-5. Members/Macros Needed/Updated by AUMPGEN Code

7.3.4 - Application Name Derivation Routine (AUMAURT)

The Analyzer's application name derivation routine assigns a value to the Application Name data element, AUMAPU. By default, AUMAPU is set to the application name defined in your MeasureWare parm file. Most installations do not need to code AUMAURT.

This user-written SAS routine is stored in prefix.MICS.PARMS(AUMAURT) and is invoked for each application processed in the AUMAPP file. The AUMAPU element in the AUMKPS and AUMIPS files obtains its value from the AUMAPP file by matching application numbers (AUMAPNUM).

You are responsible for ensuring the accuracy of the exit. A worksheet for coding the UNIX Application Name Derivation Routine is shown in Figure 7-6. For example, you can use the AUMAPU data element to make your application names more explicit for reporting purposes.

Coding Considerations

Follow these guidelines for coding this CA MICS exit routine:

1. Ensure that all of the fields that you require are available for application name construction.
2. Refer to the discussion of exit coding in Section 2.3.1.5, Notes on Coding CA MICS Parameters, in the CA MICS Planning, Installation, Operation, and Maintenance Guide.

```
+-------------------------------------------------------------------------+
| INSTALLATION PREPARATION WORKSHEET:  AUM Application Name Derivation     |
|                                      Routine Definition                 |
| PARMS Library Member is AUMAURT                                         |
| Reference Sections:  7.3.4, CA MICS Analyzer for MeasureWare Guide       |
+-------------------------------------------------------------------------+
|                                                                         |
| * VALIDATE FOR VALID APPLICATION UNITS, WHERE POSSIBLE ;                |
|     IF application data is not valid GOTO AURTOVHD ;                    |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| * BUILD APPL. UNIT FIELDS;                                              |
|     AUMAPU=field source 1 ||                                            |
|             field source n ;                                            |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
|     GOTO AURTRTEX ;                                                     |
| * LINKED ROUTINE TO BUILD INSTALLATION OVERHEAD APPLICATION UNITS;      |
| AURTOVHD:                                                               |
|     AUMAPU='overhead category' ;                                        |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| _______________________________________________________________________ |
| AURTRTEX:                                                               |
+-------------------------------------------------------------------------+
| ....5...10...15...20...25...30...35...40...45...50...55...60...65...70.. |
+-------------------------------------------------------------------------+
```

Figure 7-6. Application Name Derivation Routine Worksheet

7.3.5 - Database Space Modeling (DBMODEL)

This section describes the information required to define the Analyzer to the Database Space Modeling Facility. Specifically, the user must provide values for the cycle (data retention) definitions. The special considerations for defining the FILE input statements to the Database Space Modeling facility are discussed.

Section 2.3.4 of the PIOM, Database Space Modeling Facility, explains how these values are used in estimating the DASD requirements of the database, and how the user can input these values to the modeling facility.

This section covers the following topics:

1 - Data Retention Specifications (FILE Statements)
2 - DBMODEL Input Statements

7.3.5.1 - Data Retention Specifications (FILE Statements)

Figure 7-7 provides a worksheet to define the different retention limits for the online and archive database files. The numbers shown on the worksheet reflect the default retention values in prefix.MICS.PARMS(DBMODEL).

Use the worksheet to define your retention specifications, which, in turn, will allow you to determine the appropriate values for the modeling process.

Note that the DBMODEL worksheet shown here contains values for this product only. See Section 2.3.4.1, Preparing the Modeling Input, in the PIOM for additional information.

The worksheet is organized by information area. Each file in the area is listed by name. For each file, a line is formatted to allow six definitions in the online database and two in the archive database:

- The online database files quantify the number of cycles of data that will be maintained in the DETAIL, DAYS, WEEKS, MONTHS, and YEARS timespans and the TABLES data area.

- The two definitions for the archive database files quantify the number of cycles of data to be retained, up to the cutoff limit defined. The archive definitions have no impact on the size of the database and may be specified whether or not the weekly and/or monthly archive history files have actually been activated (see Section 2.3.3, CA MICS JCL Planning and Parameters, of the PIOM).

The worksheet formats provide an underscored area for the user's definition, followed by the recommended value, shown within parentheses. If the underscored area contains a value of 00, the file is not supported for the indicated timespan. To add support, you must perform database tailoring as described in Section 6.2, Tailoring the Database, of the System Modification Guide (SMG).

When specifying a retention limit, remember that the number may never be zero if the file has been defined to be active in the timespan.

```
+-------------------------------------------------------------------------+
| INSTALLATION PREPARATION WORKSHEET: Database Data Retention Definitions  |
|                                                                          |
| PARMS Library Member is DBMODEL                                          |
| Reference Sections:  7.3.5.1, CA MICS Analyzer for MeasureWare Guide     |
+------+----------------------------------------------+-------------------+
|      |           Online Database Retention          |  Archive Cut-Off  |
| File |DETAIL    DAYS   WEEKS   MONTHS  YEARS  TABLES | WEEKS    MONTHS   |
| Name | __(NA)  __(NA)  __(NA)  __(NA)  _(NA)  __(NA) |___ (NA)  ___ (NA) |
+------+----------------------------------------------+-------------------+
|AUMGPR| __(03)  __(07)  __(13)  __(06)  _(05)  00(00) |___ (053) ___ (036)|
|AUMAPP| __(10)  __(07)  __(13)  __(06)  _(05)  00(00) |___ (053) ___ (036)|
|AUMDDA| __(03)  __(07)  __(13)  __(06)  _(05)  00(00) |___ (053) ___ (036)|
|AUMNIF| __(03)  __(07)  __(13)  __(06)  _(00)  00(00) |___ (053) ___ (036)|
|AUMKPS| __(10)  __(00)  __(00)  __(06)  _(05)  00(00) |___ (000) ___ (036)|
|AUMIPS| __(10)  __(00)  __(00)  __(00)  _(00)  00(00) |___ (000) ___ (000)|
|AUMMCF| __(10)  __(14)  __(53)  __(24)  _(00)  00(00) |___ (053) ___ (036)|
|AUMVLM| __(03)  __(07)  __(13)  __(06)  _(05)  00(00) |___ (053) ___ (036)|
|AUMTTA| __(03)  __(07)  __(13)  __(06)  _(05)  00(00) |___ (053) ___ (036)|
+------+----------------------------------------------+-------------------+
| ....5...10...15...20...25...30...35...40...45...50...55...60...65...70...|
+-------------------------------------------------------------------------+
```

Figure 7-7. Database Data Retention Definitions Worksheet

7.3.5.2 - DBMODEL Input Statements

The DBMODEL member of prefix.MICS.PARMS provides the input to the Database Space Modeling Facility.

Update the DBMODEL member using the information collected on the worksheet in Figure 7-7. To actually perform the space modeling, submit the jobs as described in Section 2.3.4.2 of the PIOM.

7.4 - Other Related Parameters

This section contains information about other parameters related to the operation and use of the Analyzer for MeasureWare.

Remember to update the following parameters when configuring this Analyzer:

- Computing System Parameters (SYSID)
- Performance Group Names (PRFGP)

Computing System Parameters (SYSID)

Data in the CA MICS database is associated with and summarized according to the computing system that produced the raw data processed by CA MICS. Each computing system is identified by its SYSID.

Information about the SYSIDs used at your installation is coded in prefix.MICS.PARMS(SYSID). Changing this information after CA MICS is in use is a two-step process: first, change the SYSID member in prefix.MICS.PARMS; then execute the job BASPGEN in the prefix.MICS.CNTL library.

CA MICS terminates with a user abend if input data comes from a SYSID that is not defined in the PARMS library.

The following text summarizes how to code parameters for prefix.MICS.PARMS(SYSID). Section 2.3.2.2 of the CA MICS Planning, Installation, Operation, and Maintenance Guide contains more information about coding prefix.MICS.PARMS(SYSID).

- The format of the statements is free-form but positional, as follows:
  - ORGSYSID (should match AUMOPS definitions)
  - Logical SYSID (can match ORGSYSID or can be a more meaningful name)
  - CPU Model (enclosed in quotes)
  - SCP (set to UNIX)
  - RMF CPU Conversion Factor (set to 0)
  - MIPS Rate (set to 0)
  - System Name (1 to 30 character name enclosed in quotes)
- A sample entry in SYSID:

```
AUM1 AUM1 'XX/XXX' UNIX 0 0 'CHICAGO DEMO SYSTEM'
```

Performance Group Names (PRFGP)

If this Analyzer is in its own database unit or in a unit with analyzers that do not require PRFGP, you do not need to code prefix.MICS.PARMS(PRFGP). (You must leave the module, which contains comments, in the unit's library.)

If this Analyzer is in a database unit with analyzers that require PRFGP, you must ensure that the SYSIDs containing UNIX data are coded in PRFGP. Refer to Section 2.3.2 of the CA MICS Planning, Installation, Operation, and Maintenance Guide for more information about coding prefix.MICS.PARMS(PRFGP).

For example, if your prefix.MICS.PARMS(AUMOPS) OPTIONS statement is

```
OPTIONS AUM1 CHICAGO_DEMO
```

and your prefix.MICS.PARMS(SYSID) statement is

```
AUM1 AUM1 'XX/XXX' UNIX 0 0 'CHICAGO DEMO SYSTEM'
```

then your prefix.MICS.PARMS(PRFGP) should contain the line

```
AUM1
```

so this Analyzer's SYSID is included in the PRFGP table.

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 8

INSTALLATION

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 8

INSTALLATION

CHAPTER 8. INSTALLATION

Install the CA MICS Analyzer Option for MeasureWare by following the instructions in the PIOM. Section 3.8 in the PIOM provides step-by-step checklists that you must follow to ensure the successful installation of the product.

Review Chapters 6 and 7 of this guide before proceeding to the installation checklists in the PIOM. These chapters provide helpful information about the data sources and the input parameters you will need to supply for proper operation of the Analyzer Option for MeasureWare. Section 6.1 provides report files that enable you to collect and export data from MeasureWare, which you then transfer to the system running CA MICS.

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 9

PROCESSING

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 9

PROCESSING

CHAPTER 9. MW OS AGENT ANALYZER PROCESSING

Creating and maintaining a large database system such as CA MICS requires closely monitoring its periodic updating processes.

It is vital to the system's effectiveness that the database be updated daily to build the required daily, week-to-date, month-to-date, and year-to-date files accurately. This data collection process for the MW OS Agent Analyzer is vital to the proper operation of this product.

The processing flow of the MW OS Agent Analyzer is described in the following sections:

1 - Processing Considerations
2 - Daily Update Processing Flow
3 - MICSLOG Excerpt from DAY089
4 - Change the Number of Work Files for DAY089
5 - Enable Internal Step Restart

9.1 - Processing Considerations

The DAY089 step of the DAILY update job is the step that processes MeasureWare data and presents it to the MW OS Agent Analyzer for processing.

Section 6.1 of this guide describes how to obtain data for processing. Before CA MICS can process the data, it must be transported to the MVS system on which CA MICS is running, meeting the requirements noted in that section's "Transferring Data" discussion.

9.2 - Daily Update Processing Flow

The structure of the DAY089 step of the CA MICS DAILY job is illustrated below:

Module ...	Calls module ...	Description	Created by ...
INCLLIB(DAY089)	USOURCE($BASMSTR)	BAS master control macros	BASPGEN
	SOURCE(CHECK,	checkpoint validation code	delivered code
	$BASSFD,	BAS component SFD flags	BASCGEN
	$AUMSFD)	AUM component SFD flags	AUMCGEN
	USOURCE(#BASEXIT)	BAS unit level exit overrides	installation tailored
	USOURCE($AUMCYCS,	AUM unit cycle definition	CYCLEGEN
	$AUMMSTR,	AUM master processing code	AUMPGEN
	$AUMEXIT)	AUM unit level exit overrides	installation tailored
	INCLLIB(BASACT)	CA MICS Accounting and Chargeback code	delivered code
	USOURCE(#USRMAC)	installation-defined macros	installation tailored
	SOURCE(DYAUMFM1,	AUM DAILY format routine	delivered code
	DYAUMFM2,	AUM DAILY format routine	delivered code
	DYAUMSUM,	AUM summarization routines	AUMCGEN
	DYAUMAGE,	AUM file aging routines	CYCLEGEN
	POST)	checkpoint update routine	delivered code

The functions performed during the DAY089 step are:

- o Input Raw Data
- o Information Area Processing
- o Database TimeSpan Processing
- o File Aging

and are described below. This discussion presumes that data has been moved to the MVS system following the guidance provided in Section 6.1 of this guide.

Input Raw Data

In this phase, the exported data is read and organized by an INFILE exit that is called by the %AUMMWAIN macro in prefix.MICS.USER.SOURCE($AUMMSTR).

The INFILE exit, which resides in sharedprefix.MICS.LOAD(AUMAIFUE), inspects the header records provided in the exported data, requests the SAS variable locations for the metrics that are collected and have an associated CA MICS element name in sharedprefix.MICS.SOURCE(AUMMWAV), and stores the information internally.

User Exits: during this phase, the USRSEL exit is invoked. This exit is described in Chapter 10 of this guide.

Information Area Processing

The DYAUMFM1 and DYAUMFM2 modules in sharedprefix.MICS.SOURCE use the internal information stored by the INFILE exit to create the DETAIL timespan files.

The nature of the MeasureWare data requires that the Analyzer reverse-calculate some elements at the DETAIL timespan. This occurs because the MeasureWare data is less granular than CA MICS needs it to be. The CA MICS data dictionary entries for such elements will show that, at the DETAIL timespan, the element is derived from raw MeasureWare metrics while at the higher timespans, the elements are derived from CA MICS elements. For example, GPRCPQTM (CPU Queue Wait Time) at the DETAIL timespan is calculated by CPU_QUEUE*INTERVAL.

User Exits: during this phase, the USRSfff (file selection) and USRAfff (accounting) exits are invoked. These exits are described in Chapter 10 of this guide.

Database Timespan Processing

Using the newly created DETAIL files, the DAYS, WEEKS, and MONTHS timespans are updated to retain the latest day, week-to-date, and month-to-date data.

There are no exits provided in this phase.

File Aging

The files are aged to delete the oldest cycle and rename the work files to be the current cycle.

There are no exits provided in this phase.

9.3 - MICSLOG Excerpt from DAY089

Messages are written to the MICSLOG during DAY089 processing. An extract from the MICSLOG follows. As you review it, notice that the AUM00064I message can include negative values.

Negative values occur for the following reasons:

- The record is either a title, header, or configuration data record because they contain invalid time stamps
- Unrecognized records are input
- Records are dropped by the USRSEL exit
- Records are dropped because VERIFY EXCLUDE is enabled in prefix.MICS.PARMS(AUMOPS) and the UNIX system identifier does not match with any of those defined in AUMOPS.

Under these circumstances, the record is not processed through the CA MICS checkpoint processor.

To ensure that the MW OS Agent Analyzer's total record input count can be reconciled with the CA MICS Platform's BAS00326, BAS00331, and BAS00337 message record counts, the MW OS Agent Analyzer subtracts (using the negative values) the unprocessed records from its total record count. The result, called the Net Total, matches the record counts in the BAS messages.

```
Tuesday, 2 January 1996------------------------------------- MICSLOG ----------------------------------------
15.28.11 BAS00119I AT CHECKPOINT FOR STEP "DAY089" IN JOB "MICSPROD" FOR DATA BASE "UNIX".
15.28.11 BAS00122I REQUIRED PREDECESSOR STEP IS:
15.28.11 BAS00122I          PROCESS=DAY
15.28.11 BAS00122I          STEP   =088
15.28.11 BAS00125I CHECKPOINT STEPS: D_STEP=088 W_STEP=900 M_STEP=900 Y_STEP=900 B_STEP=900 R_STEP=900
15.28.11 BAS00125I

15.29.40 BAS00318I ORGSYSID/COMPONENT CHECKPOINT INFORMATION BEFORE PROCESSING
15.29.40 BAS00318I

15.29.41 AUM00051I START OF INPUT FROM AUMMWA02 = MWA.HP10.D951214.DATA

15.29.41 AUM00054I    TITLE: @MICS-AUM0000-HPUX10  Export 12/14/95 Logfile: /var/opt/perf/datafiles/logglob
15.29.41 AUM00054I    ORGSYSID: CS02    HOST NAME: lgntcs2
15.29.41 AUM00054I    COLLECTOR: SCOPE/UX B.02.00    EXPORT TIME:  14DEC1995:15:55:49
15.29.41 AUM00055I    DATA TYPE: GLOB    TIME:  14DEC1995:10:00:03
15.29.41 AUM00055I    DATA TYPE: APPL    TIME:  14DEC1995:10:00:03
15.29.41 AUM00055I    DATA TYPE: PROC    TIME:  14DEC1995:10:00:03
15.29.45 AUM00055I    DATA TYPE: CONF    TIME:  13DEC1995:14:34:30
15.29.45 AUM00055I    DATA TYPE: DISK    TIME:  14DEC1995:10:00:03
15.29.45 AUM00055I    DATA TYPE: LVOL    TIME:  14DEC1995:10:00:03
15.29.46 AUM00055I    DATA TYPE: NETS    TIME:  14DEC1995:10:00:03

15.29.46 AUM00063I MWA/HPUX     DATA RECORDS READ FROM DDNAME: AUMMWA02
15.29.46 AUM00064I
15.29.46 AUM00064I    HP MeasureWare Data Type                     Count
15.29.46 AUM00064I    ----------------------------------------     -----
15.29.46 AUM00064I    Configuration
15.29.46 AUM00064I         Title ............................          1
15.29.46 AUM00064I         Header 1 .........................          1
15.29.46 AUM00064I         Header 2 .........................          1
15.29.46 AUM00064I         Data .............................          3
15.29.46 AUM00064I         Total Records ....................          6
15.29.46 AUM00064I
15.29.46 AUM00064I    Global
15.29.46 AUM00064I         Title ............................          1
15.29.46 AUM00064I         Header 1 .........................          1
15.29.46 AUM00064I         Header 2 .........................          1
15.29.46 AUM00064I         Data .............................          6
15.29.46 AUM00064I         Total Records ....................          9
15.29.46 AUM00064I
```

Figure 9-1. MICSLOG Extract from DAY089 for an HP-UX System at B.10 (Part 1 of 3)

```
15.29.46 AUM00064I    Application
15.29.46 AUM00064I         Title ............................          1
15.29.46 AUM00064I         Header 1 ........................          1
15.29.46 AUM00064I         Header 2 ........................          1
15.29.46 AUM00064I         Data ............................         18
15.29.46 AUM00064I         Total Records ...................         21
15.29.46 AUM00064I
15.29.46 AUM00064I    Process
15.29.46 AUM00064I         Title ............................          1
15.29.46 AUM00064I         Header 1 ........................          1
15.29.46 AUM00064I         Header 2 ........................          1
15.29.46 AUM00064I         Data ............................        765
15.29.46 AUM00064I         Total Records ...................        768
15.29.46 AUM00064I
15.29.46 AUM00064I    Network Interface
15.29.46 AUM00064I         Title ............................          1
15.29.46 AUM00064I         Header 1 ........................          1
15.29.46 AUM00064I         Header 2 ........................          1
15.29.46 AUM00064I         Data ............................          6
15.29.46 AUM00064I         Total Records ...................          9
15.29.46 AUM00064I
15.29.46 AUM00064I    Disk
15.29.46 AUM00064I         Title ............................          1
15.29.46 AUM00064I         Header 1 ........................          1
15.29.46 AUM00064I         Header 2 ........................          1
15.29.46 AUM00064I         Data ............................          6
15.29.46 AUM00064I         Total Records ...................          9
15.29.46 AUM00064I
15.29.46 AUM00064I    Logical Volume
15.29.46 AUM00064I         Title ............................          1
15.29.46 AUM00064I         Header 1 ........................          1
15.29.46 AUM00064I         Header 2 ........................          1
15.29.46 AUM00064I         Data ............................         24
15.29.46 AUM00064I         Total Records ...................         27
15.29.46 AUM00064I
15.29.46 AUM00064I    Transaction
15.29.46 AUM00064I         Title ............................          0
15.29.46 AUM00064I         Header 1 ........................          0
15.29.46 AUM00064I         Header 2 ........................          0
15.29.46 AUM00064I         Data ............................          0
15.29.46 AUM00064I         Total Records ...................          0
15.29.46 AUM00064I
15.29.46 AUM00064I    Unrecognized Data Records ............          0
15.29.46 AUM00064I
15.29.46 AUM00064I    Records Dropped by _USRSEL ...........          0
15.29.46 AUM00064I
15.29.46 AUM00064I    Records Dropped for VERIFY EXCLUDE ....          0
15.29.46 AUM00064I
```

Figure 9-1. MICSLOG Extract from DAY089 for an HP-UX System at B.10 (Part 2 of 3)

```
15.29.46 AUM00064I    Total
15.29.46 AUM00064I         Overall ..........................        849
15.29.46 AUM00064I        (Title and Header Records) ........        -21
15.29.46 AUM00064I        (Configuration Data Records) ......         -3
15.29.46 AUM00064I        (Unrecognized Records) ............          0
15.29.46 AUM00064I        (Records Dropped by _USRSEL) ......          0
15.29.46 AUM00064I        (Records Dropped - VERIFY EXCLUDE).          0
15.29.46 AUM00064I         Net Total .......................        825
15.29.46 AUM00064I
15.29.49 AUM00064I

15.29.49 AUM00060I DYAUMFMT INPUT PROCESS HAS COMPLETED.
15.29.49 BAS00326I INPUT FILE(S) CONTAINED 1566 RECORDS FROM 07DEC1995:13:01:01.00 THRU 14DEC1995:10:29:03.00
15.29.49 BAS00326I
15.29.49 BAS00326I
15.29.49 BAS00331I DATA SELECTION STATISTICS BY SYSTEM/COMPONENT
15.29.49 BAS00331I
15.29.49 BAS00331I                                  - - - - - - - - - - - DELETED - - - -
15.29.49 BAS00331I   SYSTEM COMPONENT        INPUT    DATE RANGE     SHORT RECORDS   OPTION       KEPT
15.29.49 BAS00331I
15.29.49 BAS00343I    CS02  AUM                  825             0              0             0            825

15.29.49 BAS00337I   TOTAL RECORDS KEPT = 825
```

Figure 9-1. MICSLOG Extract from DAY089 for an HP-UX System at B.10 (Part 3 of 3)

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 10

MODIFICATION

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

CHAPTER 10

MODIFICATION

Chapter 10. MODIFICATION

The CA MICS Analyzer for MeasureWare is tailored primarily using options and parameters and secondarily using exits. Most of the customization is done by following the standard procedures defined by the CA MICS System Modification Guide (SMG). This chapter describes the Analyzer's exits that are available during daily update processing.

The following section is provided:

1 - Standard User Exits

10.1 - Standard User Exits

Before implementing an exit, determine whether the required modification can be accomplished by parameter modification or file tailoring.

At times, additional system customization is needed to tailor CA MICS to meet your installation's needs. There are many user exit points provided within CA MICS. Each of these points corresponds to a position in the CA MICS logic where user modification:

- o Is likely to occur, based on our experience
- o Presents only a slight possibility of compromising the CA MICS database or operational integrity
- o Is relatively easy to document and understand

Generally, you have four methods with which to augment, change, and enhance CA MICS.

- o CA MICS standard option and parameter definitions

  CA MICS complex- and unit-level parameters provide most users with adequate provisions for customizing CA MICS.

- o Modification through user exits

  If you have installation-dependent requirements that are not adequately addressed through the standard options and parameters, one or more of the user exits provided with CA MICS may be used to insert user-written routines that satisfy the requirements.

  Using the CA MICS exit facilities to augment CA MICS processing logic is the safest method of system modification.

- o Extension through user-written components

  Although most users do not initially consider this a way of modifying a standard CA MICS product, developing a user-written component may be prudent. For example, it may be better to write a user component to handle additional data than to attempt to modify the logic of the CA MICS Analyzer for MeasureWare.

- o System code modification

  Finally, if all else fails, you may be able to satisfy your installation's requirements by implementing a source

code change.

Source code changes should only be implemented as a last resort, and after having discussed them with CA MICS Product Support.

General Exit Considerations

You should carefully explore whether or not it is necessary to develop a user exit routine. You can facilitate this process by discussing the requirement with CA MICS Product Support.

Designing, coding, testing, and implementing CA MICS user exit routines should be done carefully, since any resulting errors may corrupt the data. You should be meticulous in defining and validating exit routines to ensure that system integrity and performance are not adversely affected.

If an exit must be used, refer to Chapter 9 of this guide for a description of the location of the exit in the DAILY update. For a better understanding of CA MICS user exits, see Section 4.3 of the System Modification Guide.

Types of Exits

There are two types of user exits that you can use to modify CA MICS logic: general exits and product exits.

The general exits are used by all installed products. They allow the insertion of user exit routines that extend CA MICS logic.

Product exits apply only to the product for which they are provided. They are classified as to their purpose and are comprised of input, output, parameter-related, and special exits.

- Input exits are used during the input phase of the product update process. They are invoked as the raw input records are read.

  This Analyzer provides the input exit _USRSEL. You can use this exit to select or exclude input data that meets special installation-defined criteria, add data elements, or modify record fields as they are read.

- Parameter-related exits are defined as part of the standard CA MICS installation process and are normally associated with other parameters or options.

  AUMACRT and AUMAURT are parameter-driven exits that are identified in this chapter and documented in Chapter 7 of this guide.

o Output exits are used when records are being written to the CA MICS database. They are invoked just before CA MICS writes an observation to the CA MICS database.

This Analyzer provides one output exit routine for each product file. You can use this output exit routine to selectively block the writing of specified records, alter data elements in the records, or produce additional records for the target file or a user-defined file.

The user exits for the CA MICS Analyzer for MeasureWare are identified below.

General Exits

USRIHL - Checkpoint File Inspection

This exit is documented in Section 4.3.2.1 of the CA MICS System Modification Guide.

Product Input Exits

USRSEL - File Processing Selection

Product Parameter-Related Exits

AUMACRT - Account Code Derivation Routine
AUMAURT - Application Name Derivation Routine

Product Output Exits

USRSGPR - UNIX Global Performance File Exit
USRSAPP - UNIX Application Performance File Exit
USRSDDA - UNIX Disk Activity File Exit
USRSNIF - UNIX LAN Activity File Exit
USRSKPS - UNIX Process Activity File Exit
USRSIPS - UNIX Process Exception File Exit
USRSMCF - UNIX System & Collector Configuration Profile Exit
USRSVLM - UNIX Logical Volume Activity File Exit
USRSTTA - UNIX Business Transaction SLO Activity File Exit

The following sections provide detailed information:

1 - Input Exits
2 - Parameter-Related Exits
3 - Output Exits

10.1.1 - Input Exits

This section describes the USRSEL exit, which is shared by the MW OS Agent Analyzer with the other installed CA MICS products.

USRSEL	Input Record Selection Exit (USRSEL)

DESCRIPTION: The _USRSEL exit enables you to reject records from processing during the Analyzer's input format routine (DYAUMFMT). It can be defined at the complex level in sharedprefix.MICS.SOURCE(#BASEXIT) or at the unit level in prefix.MICS.SOURCE(#BASEXIT).

INVOCATION: For this Analyzer, the _USRSEL exit gains control during the input format routine DYAUMFMT in the DAY089 step of the CA MICS DAILY job.

ACCOUNTING INTERFACE: No interface is provided.

USES: Use this exit to exclude records from processing during the Analyzer's daily input step. To exclude a record, set SKIP_REC to one. For example, you can code the following to skip records that contain the system ID "TEST":

```
MACRO _USRSEL;
  IF ROUTINE EQ 'DYAUMFMT' THEN DO;
   IF ORGSYSID EQ 'TEST' THEN SKIP_REC=1;
  END;
%
```

The Analyzer for MeasureWare provides the appropriate logic to delete the record based on the value of SKIP_REC after the _USRSEL exit is invoked.

ELEMENTS AVAILABLE: ORGSYSID - Original System ID

CODING RESTRICTIONS: For additional coding restrictions, refer to the System Modification Guide, Section 4.3.2.1.

10.1.2 - Parameter-Related Exits

This section identifies the parameter-related exits you define at product installation. Detailed descriptions of these exits are presented in Chapter 7 of this guide.

AUMACRT — Account Code Derivation Routine

DESCRIPTION: The AUMACRT exit derives the value of the account code elements defined in sharedprefix.MICS.PARMS(AUMACCT). It is a complex-level exit, so its definition applies to all database units in which the Analyzer for MeasureWare is installed.

INVOCATION: The AUMACRT exit gains control during the input format routine (DYAUMFMT) of the daily update processing step DAY089. It is invoked once for each killed process record in the AUMKPS file.

ACCOUNTING INTERFACE: No interface is provided.

USES: This exit sets the value of the AUMACTx variables.

ELEMENTS AVAILABLE: All DETAIL timespan elements of the AUMKPS file are available.

CODING RESTRICTIONS: Refer to Section 7.2.3 of this guide.

SPECIAL NOTES: The exit code is contained in sharedprefix.MICS.PARMS(AUMACRT).

SAMPLE USER EXIT: See Section 7.1 for a sample exit.

AUMAURT — Application Name Derivation Exit

DESCRIPTION: The AUMAURT exit derives the value of the UNIX application name (AUMAPU). Because MWA data provides the application name, this exit does not normally need to be coded.

INVOCATION: The AUMAURT exit gains control during the input format routine (DYAUMFMT) of the daily update processing step DAY089. It is invoked once for each application data record processed by the Analyzer for MeasureWare.

ACCOUNTING INTERFACE: No interface is provided.

USES: This exit sets the value of the AUMAPU data element.

ELEMENTS AVAILABLE: All DETAIL timespan elements of the files are available.

CODING RESTRICTIONS: Refer to Section 7.3.4 of this guide.

SPECIAL NOTES: The exit code is contained in sharedprefix.MICS.PARMS(AUMAURT).

10.1.3 - Output Exits

Output exits are invoked just prior to adding the observation to the DETAIL timespan of the file. You can use these exits to modify the value of elements or to prevent selected observations from being added.

USRSfff	UNIX File Selection Exit

DESCRIPTION: The USRSfff exit enables you to modify or select observations immediately prior to output. The files available (fff) are GPR, APP, KPS, IPS, NIF, DDA, MCF, VLM, and TTA.

INVOCATION: These exits are invoked in the information area processing phase of the DYAUMFMT routine, immediately before the file is written.

ACCOUNTING INTERFACE: No interface is provided.

USES: The exit allows elements to be modified and observations to be selected from processing.

ELEMENTS AVAILABLE: All elements in the file are available.

CODING RESTRICTIONS: Refer to the System Modification Guide, Section 4.3.1.2.

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

APPENDIX A

MESSAGES

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

APPENDIX A

MESSAGES

Appendix A. MESSAGES

This appendix documents the messages generated by the CA MICS Analyzer Option for MeasureWare.

Message Number Range	Produced During	Written to
AUM00001 - AUM01999	Operational processing	MICSLOG
AUM02000 - AUM02999	MICF reporting	MICSLOG
AUM03000 - AUM03999	Infile exit processing	SAS log

The messages are listed in ascending numerical sequence and include the full text of the message, the type (explained in the following table), the reason for the message, appropriate user action, and applicable documentation references.

Message Type	Explanation
Information	An Analyzer option or potentially important feature in the data.
Warning	A condition of either the data or the control statements that does not affect the Analyzer's operation but may lead to unexpected results.
Error	A problem encountered with a control statement that prevents the successful completion of the Analyzer's processing. Execution stopped after all control statements are processed.

AUM00001

TEXT: AUMPGEN INPUT: DSN=%PDSMEM

TYPE: Information

REASON: An input data set is processed.

ACTION: None required.

REFERENCES: None

AUM00002

TEXT: >ERR> ORGSYSID GREATER THAN 4 CHARACTERS IN LENGTH: %KEYWORD2.

TYPE: Error

REASON: The value coded for the original SYSID in AUMOPS must be 4 or fewer characters.

ACTION: Correct the ORGSYSID on the OPTIONS statement in prefix.MICS.PARMS(AUMOPS) and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

AUM00003

TEXT: >ERR> REQUIRED PARAMETERS NOT FOUND ON %KEYWORD1 STATEMENT.

TYPE: Error

REASON: A parameter is missing from the statement %KEYWORD1.

ACTION: Review the parameters listed for the %KEYWORD1 statement and ensure that all of the required parameters are coded in prefix.MICS.PARMS(AUMOPS). Rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

AUM00004

TEXT: >ERR> REQUIRED UNIX SYSTEM NAME NOT FOUND ON OPTIONS STATEMENT.

TYPE: Error

REASON: The OPTIONS statement must include the UNIX system name.

ACTION: Code the UNIX system name on the OPTIONS statement in prefix.MICS.PARMS(AUMOPS) and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00005  |
+------------+

   TEXT:  SCP SPECIFIED IN PARMS(SYSID) FOR ORGSYSID
            %KEYWORD2
          IS NOT A UNIX SYSTEM.  VALID VALUE IS UNIX.

   TYPE:  Warning

   REASON:  prefix.MICS.PARMS(SYSID) SCP keyword must be UNIX
            for this unit.

   ACTION:  Change the SCP keyword in the SYSID parameter to
            UNIX and rerun BASPGEN and AUMPGEN.

   REFERENCES:  Planning, Installation, Operation, and
                Maintenance Guide, Section 2.3.2.2

+------------+
|  AUM00006  |
+------------+

   TEXT:  >ERR> REQUIRED DDNAME FOR ORGSYSID %KEYWORD2 NOT
            FOUND ON %KEYWORD1 STATEMENT.

   TYPE:  Error

   REASON:  The name of the input data set is missing from
            the MWA statement in prefix.MICS.PARMS(AUMOPS).

   ACTION:  Code the correct ddname on the MWA statement and
            rerun AUMPGEN.  The ddname must also be defined
            in INPUTAUM.

   REFERENCES:  Section 7.3.1 in this guide

+------------+
|  AUM00007  |
+------------+

   TEXT:  >ERR> OPTIONS ENTRY FOR ORGSYSID %SYSNAME HAS BEEN
            DEFINED MORE THAN ONCE.

   TYPE:  Error

   REASON:  More than one OPTIONS statement for the original
            system identifier %SYSNAME was coded in AUMOPS.

   ACTION:  Correct the OPTIONS statement in AUMOPS and rerun
            AUMPGEN.

   REFERENCES:  Section 7.3.1 in this guide

+------------+
|  AUM00008  |
+------------+

   TEXT:  >ERR> SYSID %KEYWORD2 NOT DEFINED IN PARMS(SYSID)

   TYPE:  Error

   REASON:  The original system identifier specified in
            AUMOPS was not defined in
            prefix.MICS.PARMS(SYSID).

   ACTION:  Either add this system identifier to
            prefix.MICS.PARMS(SYSID), run BASPGEN, then run
            AUMPGEN; or correct the original system
            identifier in AUMOPS and run AUMPGEN.

   REFERENCES:  Sections 7.3.1 and 7.3.2 in this guide

+------------+
|  AUM00009  |
+------------+

   TEXT:  >ERR> ORGSYSID %KEYWORD2 ON %KEYWORD1 STATEMENT
          IS NOT THE SAME AS THE ONE ON THE ADJACENT
            OPTIONS STATEMENT.

   TYPE:  Error

   REASON:  The original system identifier specified on the
            adjacent OPTIONS and MWA statements in AUMOPS
            must be the same.

   ACTION:  Change the original system identifier on the
            OPTIONS or MWA statements and rerun AUMPGEN.

   REFERENCES:  Section 7.3.1 in this guide

+------------+
|  AUM00010  |
+------------+

   TEXT:  >ERR> INVALID KEYWORD %KEYWORD4 ON %KEYWORD1
            STATEMENT.
          VALID KEYWORD: FORCE <OPTIONS> OR ERRSKIP

   TYPE:  Error

   REASON:  A value other than FORCE or ERRSKIP was specified
```

on the %KEYWORD1 statement.

ACTION: Correct the value on the %KEYWORD1 statement and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00011  |
+------------+
```

TEXT: >ERR> PREVIOUS OPTIONS ENTRY IS A DANGLING STATEMENT.
NO MWA STATEMENT ASSOCIATES WITH IT.

TYPE: Error

REASON: A system is defined on an AUMOPS OPTIONS statement, but that system is not associated with MWA data in a MWA/datasource statement.

ACTION: Either delete the extraneous OPTIONS statement or add an MWA/datasource statement using the system. Then rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00012  |
+------------+
```

TEXT: >ERR> SUBPARAMETER "%KEYWORD2" IS NOT VALID ON A %KEYWORD1 STATEMENT.

TYPE: Error

REASON: A value other than ABORT, NOABORT, or EXCLUDE was specified on the %KEYWORD1 statement.

ACTION: Correct the value on the %KEYWORD1 statement and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00013  |
+------------+
```

TEXT: >ERR> MORE THAN ONE MWA STATEMENT DEFINED.
OPTIONS AND MWA STATEMENTS MUST BE USED IN PAIRS.

TYPE: Error

REASON: Each OPTIONS statement has an MWA statement. There is an MWA statement in AUMOPS that does not have an associated OPTIONS statement.

ACTION: Either delete the extraneous MWA statement or add an OPTIONS statement for the MWA statement. Then rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00014  |
+------------+
```

TEXT: >ERR> COMMON INPUT DD FOR %DTYPE HAS BEEN DEFINED MORE THAN ONCE.

TYPE: Error

REASON: A COMMON statement for the named data source was defined more than once.

ACTION: Delete the extraneous COMMON statement and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00015  |
+------------+
```

TEXT: >ERR> INPUTAUM DDNAME IS LONGER THAN 8 CHARACTERS: %X

TYPE: Error

REASON: The ddname of this input data set in INPUTAUM is too long.

ACTION: Correct the name in INPUTAUM and rerun AUMPGEN.

REFERENCES: Section 7.3.2 in this guide

```
+------------+
|  AUM00016  |
+------------+
```

TEXT: >ERR> KEYWORD "%KEYWORD1" IS NOT RECOGNIZED - INVALID STATEMENT.

TYPE: Error

REASON: An invalid keyword, %KEYWORD1, was specified in AUMOPS.

ACTION: Verify and correct the statement specified in AUMOPS and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

AUM00017

TEXT: END OF MEMBER %PMEM INPUT.

TYPE: Information

REASON: The member, %PMEM, was processed.

ACTION: None required.

REFERENCES: None

AUM00018

TEXT: >ERR> REQUIRED STATEMENT MISSING: %STR

TYPE: Error

REASON: The statement, %STR, is missing from AUMOPS.

ACTION: Add the statement and its specifications to AUMOPS and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

AUM00019

TEXT: ERROR(S) ENCOUNTERED IN AUM UNIT PARAMETER DEFINITIONS; MACRO GENERATION ABORTED. CORRECT THE ERROR(S), AND RE-EXECUTE AUMPGEN.

TYPE: Error

REASON: Messages earlier in the MICSLOG specify the problems encountered while processing the input data.

ACTION: Correct the problems and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

AUM00020

TEXT: >ERR> DUPLICATE DDNAMES NOT ALLOWED: %DDNAME

TYPE: Error

REASON: The ddname on the MWA statement in AUMOPS was also specified on the COMMON statement or another MWA statement.

ACTION: Correct the ddname on either the MWA or COMMON statements and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

AUM00021

TEXT: >ERR> SPECIFIED DDNAME: %KEYWORD3 IS GREATER THAN 8 CHARACTERS IN LENGTH.

TYPE: Error

REASON: The ddname specified in AUMOPS must be 8 or fewer characters long.

ACTION: Change the ddname in AUMOPS, verify that it matches the ddname specified in INPUTAUM, and rerun AUMPGEN.

REFERENCES: Sections 7.3.1 and 7.3.2 in this guide

AUM00022

TEXT: >ERR> REQUIRED OPTIONS STATEMENT NOT FOUND

TYPE: Error

REASON: An OPTIONS statement was not coded prior to the MWA statement.

ACTION: Correct AUMOPS and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00023  |
+------------+
```

TEXT: >ERR> MEMBER %PMEM IS EMPTY OR CONTAINS NO DD STATEMENTS.

TYPE: Error

REASON: INPUTAUM was not defined.

ACTION: Specify a ddname in INPUTAUM and rerun AUMPGEN.

REFERENCES: Section 7.3.2 in this guide

```
+------------+
|  AUM00024  |
+------------+
```

TEXT: >ERR> DDNAME %DDNAME NOT FOUND IN INPUTAUM

TYPE: Error

REASON: The ddname specified on the MWA statement in AUMOPS does not match any ddname in INPUTAUM.

ACTION: Correct the MWA statement or INPUTAUM and rerun AUMPGEN.

REFERENCES: Sections 7.3.1 and 7.3.2 in this guide

```
+------------+
|  AUM00025  |
+------------+
```

TEXT: INPUTAUM DDNAME: %DDNAME NOT REFERENCED IN AUMOPS

TYPE: Warning

REASON: INPUTAUM includes ddnames not specified on the MWA statement in AUMOPS.

ACTION: None required. This is a consistency check only. Verify that all data you want processed by CA MICS is referenced in AUMOPS. Then, rerun AUMPGEN if changes to AUMOPS or INPUTAUM are made.

REFERENCES: Sections 7.3.1 and 7.3.2 in this guide

```
+------------+
|  AUM00026  |
+------------+
```

TEXT: >ERR> INVALID COMMON %KEYWORD2 DDNAME: %KEYWORD3

TYPE: Error

REASON: %KEYWORD3 cannot be specified as a ddname on the COMMON %KEYWORD2 statement.

ACTION: Correct the COMMON statement in AUMOPS and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00027  |
+------------+
```

TEXT: EXTRANEOUS STATEMENT: %STR

TYPE: Warning

REASON: A statement is coded in AUMOPS that is not being used because it is not referenced.

ACTION: None required. To avoid confusion later, delete the extra statements.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00029  |
+------------+
```

TEXT: >ERR> INVALID FORCE OPTION KEYWORD: %STR

TYPE: Error

REASON: The FORCE keyword in AUMOPS specifies an invalid option.

ACTION: Replace the incorrect option with a valid one.

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00030  |
+------------+
```

```
TEXT:  >ERR> MISSING VALUE ASSOCIATED WITH THE FORCE
         OPTION: %STR

TYPE:  Error

REASON:  The FORCE keyword in AUMOPS is missing an option
         that is required.

ACTION:  Add an option to the FORCE keyword.

REFERENCES:  Section 7.3.1 in this guide
```

```
+------------+
|  AUM00031  |
+------------+
```

```
TEXT:  >ERR> MISSING A FORCE OPTION AND ITS ASSOCIATED
         VALUE

TYPE:  Error

REASON:  The FORCE keyword is specified in AUMOPS, but a
         required option is missing.

ACTION:  Add an option to the FORCE keyword.

REFERENCES:  Section 7.3.1 in this guide
```

```
+------------+
|  AUM00032  |
+------------+
```

```
TEXT:  >ERR> MISSING A CLOSING QUOTATION MARK

TYPE:  Error

REASON:  The hostname on the OPTIONS statement in
         prefix.MICS.PARMS(AUMOPS) is missing a closing
         quotation mark.

ACTION:  Correct AUMOPS and rerun AUMPGEN.

REFERENCES:  Section 7.3.1 in this guide
```

```
+------------+
|  AUM00033  |
+------------+
```

```
TEXT:  >ERR> UNRECOGNIZED KEYWORD: %STR

TYPE:  Error

REASON:  The keyword %STR follows the hostname on the
         OPTIONS statement in AUMOPS.  Additional keywords
         are not supported.

ACTION:  Remove the text following the hostname on the
         OPTIONS statement and rerun AUMPGEN.

REFERENCES:  Section 7.3.1 in this guide
```

```
+------------+
|  AUM00034  |
+------------+
```

```
TEXT:  >ERR> SPECIFIED UNIX SYSTEM ID IS LONGER THAN 40
         CHARACTERS

TYPE:  Error

REASON:  The hostname on the OPTIONS statement in AUMOPS
         can only be 40 characters long.

ACTION:  Correct the hostname in AUMOPS and rerun AUMPGEN.

REFERENCES:  Section 7.3.1 in this guide
```

```
+------------+
|  AUM00035  |
+------------+
```

```
TEXT:  >ERR> Only one GLOBAL statement allowed
       Two or more GLOBAL statements detected'

TYPE:  Error

REASON:  The GLOBAL FORCE statement can only be specified
         once in AUMOPS.

ACTION:  Retain one GLOBAL FORCE statement in AUMOPS and
         rerun AUMPGEN.

REFERENCES:  Section 7.3.1 in this guide
```

```
+------------+
|  AUM00036  |
+------------+
```

```
TEXT:  >ERR> The FORCE keyword must immediately follow the
         GLOBAL statement

TYPE:  Error

REASON:  You must specify the FORCE keyword after the
```

```
             GLOBAL statement in AUMOPS.

    ACTION:  After the GLOBAL statement in AUMOPS, specify the
             FORCE keyword, and rerun AUMPGEN.

    REFERENCES:  Section 7.3.1 in this guide

 +------------+
 |  AUM00037  |
 +------------+

    TEXT:  >ERR> GLOBAL FORCE must be followed by AIX, HPUX,
             SOLARIS, and/or ALL
           First 8 characters of argument coded as: %GFARG

    TYPE:  Error

    REASON:  You have specified an unsupported argument for
             the GLOBAL FORCE statement in AUMOPS.  Supported
             arguments are AIX, HPUX, SOLARIS, and/or ALL.
             The first 8 characters of the invalid argument
             are output.

    ACTION:  Specify a valid argument for the GLOBAL FORCE
             statement in AUMOPS and rerun AUMPGEN.

    REFERENCES:  Section 7.3.1 in this guide

 +------------+
 |  AUM00038  |
 +------------+

    TEXT:  >ERR> GLOBAL FORCE detected with no arguments
           Must be followed by AIX, HPUX, SOLARIS, and/or ALL

    TYPE:  Error

    REASON:  You have specified GLOBAL FORCE, but have not
             selected any supporting arguments.

    ACTION:  Specify one or more valid arguments after GLOBAL
             FORCE in AUMOPS and rerun AUMPGEN.

    REFERENCES:  Section 7.3.1 in this guide

 +------------+
 |  AUM00039  |
 +------------+

    TEXT:  >ERR> GLOBAL record not output from parsing
           Please contact CA MICS Product Support

    TYPE:  Error

    REASON:  A GLOBAL record has not been output.

    ACTION:  GLOBAL record has not been processed.

    REFERENCES:  Contact CA Technical Support.

 +------------+
 |  AUM00051  |
 +------------+

    TEXT:  START OF INPUT FROM %DDNAME = %DSNAME

    TYPE:  Information

    REASON:  Record the name of the input data set.

    ACTION:  None required.

    REFERENCES:  None

 +------------+
 |  AUM00052  |
 +------------+

    TEXT:  >ERR> INPUT PROCESSING ERROR FOR %UNXDTYPE
                 %INFXIT RETURN CODE = %ERR_CODE
           PLEASE SEE SASLOG FOR DETAIL MESSAGE

    TYPE:  Error

    REASON:  An error was detected in the infile exit.

    ACTION:  Messages in the SAS log specify problems
             encountered during processing.

    REFERENCES:  Chapter 6 in this guide

 +------------+
 |  AUM00053  |
 +------------+

    TEXT:  >ERR> INCORRECT %UNXDTYPE DATA FORMAT
           CURRENT INPUT DATA DOES NOT START WITH TITLE
             RECORD
           FIRST RECORD ID: %MWA_RID

    TYPE:  Error

    REASON:  The input data must begin with a title record.
             Processing is aborted.
```

ACTION: Review the logs of data collection and transfer to ensure completeness. Correct errors and restart the DAY089 step.

REFERENCES: Chapter 6 in this guide

```
+------------+
|  AUM00054  |
+------------+
```

TEXT: TITLE: %WRK_TTL
ORGSYSID: %ORGSYSID HOST NAME: %AUMSYSID
COLLECTOR: %COLLECTR EXPORT TIME: %TS

TYPE: Information

REASON: Record the characteristics of the data being processed. This message is issued for each title record encountered.

ACTION: None required.

REFERENCES: None

```
+------------+
|  AUM00055  |
+------------+
```

TEXT: DATA TYPE: %REC_ID TIME: %TS

TYPE: Information

REASON: Record the type and time of the first data record immediately after the title record being processed.

ACTION: None required.

REFERENCES: None

```
+------------+
|  AUM00056  |
+------------+
```

TEXT: UNSUPPORTED RECORD ID: %MWA_RID (%REC_ID)

TYPE: Warning

REASON: An unsupported record has been input.

ACTION: The data from the unsupported record is not processed.

REFERENCES: None

```
+------------+
|  AUM00057  |
+------------+
```

TEXT: INFILE EXIT RETURNED AN INVALID RECORD TYPE: %REC_ID (%MWA_RID)

TYPE: Warning

REASON: Data from an unsupported record type was presented for processing.

ACTION: The data is not processed.

REFERENCES: Contact CA Technical Support.

```
+------------+
|  AUM00058  |
+------------+
```

TEXT: INPUT DATA ORIGINATED FROM INVALID OPERATING SYSTEM: %AUMOPSYS
FORCE OPTION SPECIFIED IN AUMOPS. PROCESSING WILL CONTINUE...

TYPE: Warning

REASON: %AUMOPSYS is an unsupported operating system.

ACTION: Verify the name of the UNIX system from which the data came. DYAUMFMT will continue processing, because FORCE is specified in prefix.MICS.PARMS(AUMOPS).

REFERENCES: Section 7.3.1 in this guide

```
+------------+
|  AUM00059  |
+------------+
```

TEXT: >ERR> Unix system name found in MWA data is %UNXSYSID This name does not match the system name(s) specified in AUMOPS for DDNAME: %DDNAME
VERIFY ABORT specified in AUMOPS.

TYPE: Error

REASON: The UNIX system name specified on the OPTIONS

```
statement in AUMOPS does not match the name in
the data. DYAUMFM1 terminates because VERIFY
ABORT was specified in AUMOPS.

ACTION:  Verify the name of the UNIX system from which the
         data came; make corrections in AUMOPS and rerun
         AUMPGEN.  Then restart the DAY089 step.

REFERENCES:  Section 7.3.1 in this guide

+------------+
|  AUM00060  |
+------------+

TEXT:  %ROUTINE INPUT PROCESS HAS COMPLETED.

TYPE:  Information

REASON:  The input process, %ROUTINE, completed
         successfully.

ACTION:  None required.

REFERENCES:  None

+------------+
|  AUM00061  |
+------------+

TEXT:  %SYSID: DATA TRUNCATION DETECTED AT THE CURRENT
         INPUT DATA SET
       SHORT RECORD OCCURRED AT %MWA_ROW:%MWA_COL
       LAST GOOD RECORD TYPE: %REC_ID   LENGTH: %LRECL

TYPE:  Warning

REASON:  The input data set contains an incomplete record.

ACTION:  Check if the data set ran out of space during the
         data transfer process.  Allocate more data space
         if required.  The short record is not processed
         because it is incomplete.

REFERENCES:  None

+------------+
|  AUM00062  |
+------------+

TEXT:  UNSUPPORTED %UNXDTYPE/%UNIXOS VERSION: %SCOPEUX
       SUPPORTED VERSION(S): %SCOPELVL
       %FRCTXT OPTION SPECIFIED IN AUMOPS.  PROCESSING
         WILL CONTINUE...

TYPE:  Warning

REASON:  The version of the performance data collector is
         not supported.

ACTION:  Verify the collector's version.  Processing
         continues because FORCE or GLOBAL FORCE is
         specified in prefix.MICS.PARMS(AUMOPS).

REFERENCES:  None

+------------+
|  AUM00063  |
+------------+

TEXT:  %UNXDTYPE/%UNIXOS DATA RECORDS READ FROM DDNAME:
         %DDNAME

TYPE:  Information

REASON:  Identify the ddname that the following AUM00064
         messages are for.

ACTION:  None required.

REFERENCES:  None

+------------+
|  AUM00064  |
+------------+

TEXT:  %MLINE

TYPE:  Information

REASON:  Prints a line in the MICSLOG.  It contains the
         input record count.

ACTION:  None required.

REFERENCES:  Section 9.3 in this guide

+------------+
|  AUM00065  |
+------------+

TEXT:  NUMBER OF RECORDS DROPPED BY _USRSEL: %NSKIPS

TYPE:  Information
```

REASON: Record the number of records dropped due to the specifications in your site's USRSEL exit.

ACTION: None required.

REFERENCES: None

AUM00066

TEXT: UNIX SYSTEM NAME SPECIFIED IN PARMS(AUMOPS) DOES NOT MATCH THE ONE IN %UNXDTYPE DATA.
UNIX SYSTEM NAME IN DATA: %UNXSYSID
VERIFY NOABORT SPECIFIED IN PARMS(AUMOPS).
PROCESSING WILL CONTINUE...

TYPE: Warning

REASON: The UNIX system name specified on the OPTIONS statement in AUMOPS does not match the name in the data.

ACTION: Verify the name of the UNIX system from which the data came. Processing continues because VERIFY NOABORT was specified in AUMOPS.

REFERENCES: Section 7.3.1 in this guide

AUM00067

TEXT: >ERR> UNIX SYSTEM NAME IN %UNXDTYPE DATA DOES NOT MATCH WITH SYSTEM NAMES DEFINED IN AUMOPS
UNIX SYSTEM NAME FROM THE COMMON INPUT DD %DDNAME: %UNXSYSID

TYPE: Error

REASON: The UNIX system name in the data does not match any of the systems specified in AUMOPS for the COMMON ddname.

ACTION: Verify the name of the UNIX system from which the data came. Correct AUMOPS if necessary. CA MICS user abend 310 accompanies this message.

REFERENCES: Section 7.3.1 in this guide
Section 4.3.11 in the CA MICS Planning, Installation, Operation, and Maintenance Guide

AUM00068

TEXT: >ERR> %SYSID: MISSING CONFIGURATION DATA!
CONFIGURATION DATA ENSURES THE INTEGRITY OF THE DATA BASE
PROCESSING ABORTED.

TYPE: Error

REASON: Data collection or transfer did not include some of the configuration data required for a complete CA MICS database.

ACTION: Collect the data required for the Analyzer Option for MeasureWare processing.

REFERENCES: Chapter 6 in this guide

AUM00069

TEXT: >ERR> INPUT DATA ORIGINATED FROM INVALID OPERATING SYSTEM: %AUMOPSYS
%ROUTINE PROCESSING ABORTED.

TYPE: Error

REASON: %AUMOPSYS is an unsupported operating system.

ACTION: Verify the name of the UNIX system from which the data came. If necessary, use the FORCE option in AUMOPS to let data bypass this check.

REFERENCES: Section 7.3.1 in this guide

AUM00070

TEXT: >ERR> UNSUPPORTED %UNXDTYPE/%UNIXOS VERSION: %SCOPEUX
SUPPORTED VERSION(S): %SCOPELVL
%ROUTINE PROCESSING ABORTED.

TYPE: Error

REASON: The version of the performance data collector is not supported.

```
ACTION:  Verify the collector's version.  If necessary,
         you can bypass the version by specifying the
         FORCE VER argument in prefix.MICS.PARMS(AUMOPS).
         Use caution when specifying FORCE VER, as
         described in the referenced AUMOPS section.  If
         you change AUMOPS, rerun AUMPGEN.

REFERENCES:  Section 7.3.1 in this guide

+------------+
|  AUM00071  |
+------------+

TEXT:  UNRECOGNIZED DATA COLLECTOR: %SCOPENM
       FORCE OPTION SPECIFIED IN AUMOPS.  PROCESSING
         WILL CONTINUE...

TYPE:  Warning

REASON:  The data collector is not supported.

ACTION:  Verify the data collector name.  Processing
         continues because the FORCE option was specified
         in prefix.MICS.PARMS(AUMOPS).

REFERENCES:  Section 7.3.1 in this guide

+------------+
|  AUM00072  |
+------------+

TEXT:  >ERR> UNRECOGNIZED DATA COLLECTOR: %SCOPENM
       %ROUTINE PROCESSING ABORTED.

TYPE:  Error

REASON:  The data collector is not supported.

ACTION:  Verify the data collector name.  You can bypass
         the abend by specifying the FORCE VER argument in
         prefix.MICS.PARMS(AUMOPS).

REFERENCES:  Section 7.3.1 in this guide

+------------+
|  AUM00073  |
+------------+

TEXT:  >ERR> INPUT TO %WRKNAME FILE MAY HAVE CONTAINED
         BOTH DETAIL AND SUMMARY DATA FOR THE SAME TIME
         PERIOD
       TIME STAMP: %TS
       DURATION OF THE RECORD:  CURRENT=%DURATION
         LAST=%PREVDUR (SECONDS)
       TO PREVENT DATA BASE CORRUPTION, PROCESSING
         ABORTED.

TYPE:  Error

REASON:  Both detail data and summary data were exported
         together to the same file for the same period of
         time.

ACTION:  Choose to export either detail or summary data,
         but not both.  Define the level of data to be
         exported in the data collector's extract command
         stream.

REFERENCES:  Section 6.1 of this guide

+------------+
|  AUM00074  |
+------------+

TEXT:  NEGATIVE VALUE(S) DETECTED IN THE INPUT DATA (TIME
         = %TS)
       SEE SASLOG FOR MORE DETAILS
       AFFECTED VARIABLE(S) WILL BE SET TO SAS MISSING
         VALUE

TYPE:  Warning

REASON:  The input data from the timestamp noted contains
         negative values, which are not typical in MWA
         data.

ACTION:  Ensure that you are running the scopeux and
         extract programs at the maintenance levels noted
         in the referenced section.

REFERENCES:  Section 1.4 of this guide

+------------+
|  AUM00075  |
+------------+

TEXT:  >ERR> UNSUPPORTED %UNXDTYPE/%UNIXOS VERSION:
         %SCOPEUX
       VERSION DOES NOT MATCH WITH THE FORCED VERSION
         SPECIFIED IN AUMOPS: %FRC_VER
       SUPPORTED VERSION(S): %SCOPELVL
       %ROUTINE PROCESSING ABORTED.
```

```
TYPE:  Error

REASON:  The scopeux collector version does not match the
         version coded in prefix.MICS.PARMS(AUMOPS).

ACTION:  Verify that your scopeux collector version is
         compatible with one of the supported versions
         listed in Section 1.4 of this guide.  You can
         bypass the version by specifying the FORCE VER
         argument in prefix.MICS.PARMS(AUMOPS).  Use
         caution when specifying FORCE VER, as described
         in the referenced AUMOPS section.  If you change
         AUMOPS, rerun AUMPGEN and then restart the DAY089
         step.

REFERENCES:  Sections 1.4 and 7.3.1 of this guide

+------------+
|  AUM00076  |
+------------+

TEXT:  >ERR> UNRECOGNIZED DATA COLLECTOR: %SCOPENM
       COLLECTOR NAME DOES NOT MATCH WITH THE ONE
         SPECIFIED ON THE FORCE OPTION IN AUMOPS:
         %FRC_CLL
       %ROUTINE PROCESSING ABORTED.

TYPE:  Error

REASON:  The input data is from an unsupported version of
         the scopeux collector.  The collector version
         specified on the FORCE VER argument in
         prefix.MICS.PARMS(AUMOPS) does not match the
         collector version in the data.

ACTION:  Verify that your scopeux collector version is
         compatible with one of the supported versions
         listed in Section 1.4 of this guide.  You can
         bypass the version by specifying the FORCE VER
         argument in prefix.MICS.PARMS(AUMOPS).  Use
         caution when specifying FORCE VER, as described
         in the referenced AUMOPS section.  If you change
         AUMOPS, rerun AUMPGEN and then restart the DAY089
         step.

REFERENCES:  Sections 1.4 and 7.3.1 in this guide

+------------+
|  AUM00077  |
+------------+

TEXT:  >ERR> INPUT DATA ORIGINATED FROM INVALID OPERATING
         SYSTEM: %AUMOPSYS
       OS NAME DOES NOT MATCH WITH THE ONE SPECIFIED
         ON THE FORCE OPTION IN AUMOPS: %FRC_OS
       %ROUTINE PROCESSING ABORTED.

TYPE:  Error

REASON:  The operating system specified on the FORCE VER
         argument in prefix.MICS.PARMS(AUMOPS) does not
         match the operating system present in the data.

ACTION:  Verify that you are processing data from an
         operating system that is compatible with one of
         the supported operating systems listed in Section
         1.4 of this guide.  You can bypass the version by
         specifying the FORCE VER argument in
         prefix.MICS.PARMS(AUMOPS).  Use caution when
         specifying FORCE VER, as described in the
         referenced AUMOPS section.  If you change AUMOPS,
         rerun AUMPGEN and then restart the DAY089 step.

REFERENCES:  Sections 1.4 and 7.3.1 in this guide.

+------------+
|  AUM00078  |
+------------+

TEXT:  >CA MICS INTERNAL ERROR> UNRECOGNIZED OPERATING
         SYSTEM: %UNIXOS
       PLEASE CONTACT CA MICS PRODUCT SUPPORT

TYPE:  Error

REASON:  The operating system information from AUMPGEN
         does not correspond to an operating system that
         can be processed by the Analyzer.

ACTION:  Contact CA Technical Support.

REFERENCES:  None

+------------+
|  AUM00079  |
+------------+

TEXT:  THE FOLLOWING PROCESS RECORD WAS NOT CAPTURED IN
         APPLICATION DATA:
       SYSID: %SYSID   AUMAPNUM: %AUMAPNUM   ENDTS:
         %TS   USER: %AUMUSER   PROGRAM: %AUMPROG
       AUMAPU SET TO VALUE "OTHER"

TYPE:  Warning
```

```
REASON: There is no matching application record for the
        named process record.

ACTION:  Contact CA Technical Support.

REFERENCES: None

+------------+
|  AUM00080  |
+------------+

TEXT:  THE FOLLOWING PROCESS RECORD WAS NOT CAPTURED IN
         APPLICATION DATA:
       SYSID: %SYSID   AUMAPNUM: %AUMAPNUM   ENDTS:
         %TS   USER: %AUMUSER   PROGRAM: %AUMPROG
       AUMAPU IS SET TO VALUE MISSING

TYPE:  Warning

REASON:  There is no matching application record for the
         named process record.

ACTION:  Contact CA Technical Support.

REFERENCES:  None

+------------+
|  AUM00081  |
+------------+

TEXT:  OS NAME IN DATA (%AUMOPSYS) DOES NOT MATCH WITH
         THE MWA TYPE IN AUMOPS: %UNIXOS
       FORCE OS OPTION SPECIFIED IN AUMOPS.
         PROCESSING WILL CONTINUE...

TYPE:  Warning

REASON:  It appears that the data was not produced on the
         operating system defined in AUMOPS.

ACTION:  None required.  The FORCE VER argument was
         specified in AUMOPS, so the data is being
         processed.

REFERENCES:  Section 7.3.1 in this guide.

+------------+
|  AUM00082  |
+------------+

TEXT:  >ERR> OS NAME IN DATA (%AUMOPSYS) DOES NOT MATCH
         WITH THE MWA TYPE IN AUMOPS: %UNIXOS
       FORCE OS OPTION (%FRC_OS) DOES NOT MATCH THE OS
         NAME IN MWA DATA
       %ROUTINE PROCESSING ABORTED.

TYPE:  Error

REASON:  The operating system name specified on the FORCE
         VER argument on the MWA statement in AUMOPS does
         not match the operating system name present on
         the data.

ACTION:  Correct the FORCE VER argument and rerun AUMOPS.
         Then rerun the DAY089 step.

REFERENCES:  Sections 1.4 and 7.3.1 in this guide.

+------------+
|  AUM00083  |
+------------+

TEXT:  >ERR> OS NAME IN DATA (%AUMOPSYS) DOES NOT MATCH
         WITH THE MWA TYPE IN AUMOPS: %UNIXOS
       %ROUTINE PROCESSING ABORTED.

TYPE:  Error

REASON:  The data presented for processing did not come
         from the operating system defined on the MWA
         statement in AUMOPS.

ACTION:  Correct the OPTIONS and MWA statements in AUMOPS
         and rerun AUMPGEN. Then rerun the DAY089 step.

REFERENCES:  Sections 1.4 and 7.3.1 in this guide.

+------------+
|  AUM00084  |
+------------+

TEXT:  WILL STOP PROCESSING DATASETS AT CURRENT DD DUE TO
         A DATA PROBLEM
       DATA RECORDS INPUT FROM THE CURRENT DD WILL BE
         DROPPED
       INPUT PROCESSING WILL RESUME FROM NEXT DD DUE TO
         "ERRSKIP" SPECIFIED

TYPE:  Warning

REASON:  The data presented for processing has a problem.

ACTION:  All the data records read from the current ddname
```

are dropped from the database. DYAUMFMT will continue processing because ERRSKIP is specified in prefix.MICS.PARMS(AUMOPS). Review and correct the data problem. Then input the data to the DAY089 step again.

REFERENCES: Sections 7.3.1 in this guide.

AUM00085

TEXT: PROCESSING ABNORMALLY TERMINATED.

TYPE: Error

REASON: Input processing terminated prematurely due to an error detected by infile exit.

ACTION: Messages in the SAS log specify problems encountered during processing. Correct the problem and restart the DAY089 step.

REFERENCES: Chapter 6 in this guide.

AUM00086

TEXT: DATA RECORDS FROM THE FOLLOWING SYSID(S) WILL BE DROPPED: (DDNAME=%DDNAME)

TYPE: Warning

REASON: Input data from %DDNAME contains an error. Data records from the SYSIDs on the following AUM00087 messages are deleted from the processing because ERRSKIP is specified in prefix.MICS.PARMS(AUMOPS).

ACTION: None required.

REFERENCES: Section 7.3.1 in this guide.

AUM00087

TEXT: ---- %SYSID

TYPE: Warning

REASON: Identify the system ID whose data records are to be deleted. This message is issued with the AUM00086 message.

ACTION: None required.

REFERENCES: None

AUM00088

TEXT: >ERR> NO DATA RECORDS ARE KEPT IN THE DATA BASE
USER ABEND 998 IS BEING ISSUED

TYPE: Error

REASON: Zero records are kept in the database.

ACTION: Review and correct errors listed in the prior messages. Restart DAY089.

REFERENCES: None

AUM00089

TEXT: UNIX SYSTEM NAME SPECIFIED IN PARMS(AUMOPS) DOES NOT MATCH THE ONE IN %UNXDTYPE DATA.
UNIX SYSTEM NAME IN DATA: %UNXSYSID
DATA WILL BE EXCLUDED FROM PROCESSING DUE TO "VERIFY EXCLUDE" SPECIFIED IN AUMOPS.

TYPE: Error

REASON: The UNIX system name specified on the OPTIONS statement in AUMOPS does not match the name in the data. The input data from the noted UNIX system is discarded. DYAUMFMT processing continues because VERIFY EXCLUDE was specified in AUMOPS.

ACTION: Verify the name of the UNIX system from which the data came; make corrections in AUMOPS and rerun AUMPGEN.

REFERENCES: Section 7.3.1 in this guide

AUM00090

```
TEXT:  >CA MICS INTERNAL ERROR> UNRECOGNIZED VFYABORT
         VALUE:  %VFYABNT
       PLEASE CONTACT CA MICS PRODUCT SUPPORT

TYPE:  Error

REASON:  CA MICS internal error.

ACTION:  Contact CA Technical Support.

REFERENCES:  None
```

```
+------------+
|  AUM00091  |
+------------+

TEXT:  Obsolete %UNIXOS script detected for ORGSYSID %OR,
       GSYSID
       UPGRADE REQUIRED FOR TECHNICAL CURRENCY
       Begin using the AUMRHP12 script

TYPE:  Warning

REASON:  An old version of the script has been detected.
         In order to take advantage of technical currency,
         including the obsoletion and addition of many
         metrics, you are strongly urged to upgrade to the
         new script level detailed in this message.

ACTION:  Upgrade your scripts for all platforms and
         ORGSYSIDs that are noted in this MICSLOG message.

REFERENCES:  None
```

```
+------------+
|  AUM00092  |
+------------+

TEXT:  The GLOBAL FORCE option has been specified.
       Release validation will not occur for:

TYPE:  Informational

REASON:  The GLOBAL FORCE option has been specified in
         prefix.MICS.PARMS(AUMPGEN).  This parameter will
         bypass release validation for this product.

ACTION:  None

REFERENCES:  None
```

```
+------------+
|  AUM00093  |
+------------+

TEXT:  %GFMSGTXT

TYPE:  Informational

REASON:  No release validation will occur for HP-UX, Sun
         Solaris, and/or IBM AIX data.

ACTION:  None

REFERENCES:  None
```

```
+------------+
|  AUM00098  |
+------------+

TEXT:  >>> %PROCESS HAS COMPLETED.

TYPE:  Information

REASON:  The job, %PROCESS, completed successfully.

ACTION:  None required.

REFERENCES:  None
```

```
+------------+
|  AUM00099  |
+------------+

TEXT:  %LINE

TYPE:  Information

REASON:  Prints a line in the MICSLOG.

ACTION:  None required.

REFERENCES:  None
```

```
+------------+
|  AUM02001  |
+------------+

TEXT:  NO DATA HAS PASSED THROUGH INITIAL SELECTION.
       CHECK FILE AND SELECTION CRITERIA VALIDITY.
       RUN TERMINATES WITH A USER 998 CONDITION CODE.

TYPE:  Error
```

```
REASON:  The selection criteria specified for this MICF
         inquiry excluded all data.

ACTION:  Review data selection criteria; rerun inquiry.

REFERENCES:  None

+------------+
|  AUM02002  |
+------------+

TEXT:  A REQUIRED FILE WAS NOT FOUND IN ANY UNIT OF THE
         SPECIFIED/DEFAULT UNIT DATA BASE ID LIST.
       DBID LIST: %ICFRPDBL / RETURN CODE: %RC.
       REPORT EXECUTION IS TERMINATED.

TYPE:  Error

REASON:  The unit database from which this report is being
         generated does not contain the Analyzer file
         needed for this report.

ACTION:  Either specify a unit database that contains the
         Analyzer file and re-execute the report, or
         choose a different report.

REFERENCES:  Chapter 3 in this guide

+------------+
|  AUM02003  |
+------------+

TEXT:  THE SPECIFIED/DEFAULT UNIT DATA BASE ID LIST DOES
         NOT CONTAIN ALL REQUIRED FILES FOR THIS REPORT.
       AN ATTEMPT WILL BE MADE TO CONSTRUCT A NEW UNIT
         LIST USING ALL VALID UNITS IN THIS COMPLEX.
       REPORT EXECUTION CONTINUES.

TYPE:  Warning

REASON:  The unit database specified did not contain the
         Analyzer files needed to build this report.

ACTION:  None required.  CA MICS attempts to locate the
         required files in another unit in this complex.

REFERENCES:  None

+------------+
|  AUM02004  |
+------------+

TEXT:  UNEXPECTED ERROR DETECTED BY THE MICF DBID
         VALIDATION ROUTINE.  RETURN CODE IS %AUMWORK1.
       EXAMINE THE UNIT DATA BASE ID LIST SPECIFIED OR
         DEFAULTED FOR THIS REPORT.
       REPORT EXECUTION IS TERMINATED.

TYPE:  Error

REASON:  The MICF validation routine for database
         identifiers detected a problem with the unit
         specified for this report.

ACTION:  Review the CA MICS log and respecify the unit
         database identifier.

REFERENCES:  None

+------------+
|  AUM02005  |
+------------+

TEXT:  THE FOLLOWING UNIT DATA BASE ID LIST VALUES
         CONTAIN NONE OF THE FILES USED BY THIS REPORT:
         %MALMSUID.
       THE SPECIFIED VALUES WILL BE IGNORED.
       REPORT EXECUTION CONTINUES.

TYPE:  Warning

REASON:  The unit database you chose to use for this
         report does not contain the Analyzer files
         needed.

ACTION:  None required.

REFERENCES:  None

+------------+
|  AUM02006  |
+------------+

TEXT:  THE FOLLOWING REQUIRED FILES WERE NOT FOUND IN ANY
         OF THE DEFAULT/SPECIFIED UNITS: %MALMSFID.
       CORRECT THE UNIT DATA BASE ID LIST TO INCLUDE
         AT LEAST ONE UNIT CONTAINING THE SPECIFIED
         FILES.
       REPORT EXECUTION IS TERMINATED.

TYPE:  Error

REASON:  The unit database you chose to use for this
```

report does not contain the Analyzer files needed.

ACTION: Specify a unit that contains the required files.

REFERENCES: Chapter 3 in this guide

```
+------------+
|  AUM02007  |
+------------+
```

TEXT: NO CYCLE VALUES HAVE BEEN SPECIFIED FOR THIS REPORT.
SPECIFY AT LEAST ONE VALID CYCLE VALUE WHEN SELECTING THIS REPORT FOR EXECUTION.
REPORT EXECUTION IS TERMINATED.

TYPE: Error

REASON: You must specify the cycles of data on which you want to report.

ACTION: Specify a cycle value or range of cycles; then rerun the report.

REFERENCES: Chapter 3 in this guide

```
+------------+
|  AUM02008  |
+------------+
```

TEXT: NOTE: SOME MEASUREMENTS REPORTED BY THIS INQUIRY ARE NOT AVAILABLE IN THE FOLLOWING
OPERATING SYSTEM(S): %OS
TEXT4

TYPE: Information

REASON: The metrics supported vary by operating system. The documentation for each report lists the CA MICS elements used, and notes if an element is not available from an operating system's data.

ACTION: None required.

REFERENCES: Chapter 3 in this guide.

```
+------------+
|  AUM02009  |
+------------+
```

TEXT: NOTE: THIS REPORT IS NOT APPLICABLE TO THE
FOLLOWING OPERATING SYSTEM(S):
%OS

TYPE: Information

REASON: The metrics supported vary by operating system. The documentation for each report lists the CA MICS elements used, and notes if a report is not available from an operating system's data.

ACTION: None required.

REFERENCES: None

```
+------------+
|  AUM02010  |
+------------+
```

TEXT: THE LINES PER PAGE WAS SET TO LESS THAN 30. THIS REPORT REQUIRES A MINIMUM
PAGE SIZE OF 30. PAGE SIZE SET TO 30

TYPE: Information

REASON: To correctly format this report, the page size must be at least 30. The page size was automatically reset.

ACTION: None required.

REFERENCES: None

```
+----------+
| AUM03001 |
+----------+
```

TEXT: Unknown function request type.

TYPE: Error

REASON: The infile exit detects an unsupported function request type. Processing aborted.

ACTION: Contact CA Technical Support.

REFERENCES: None

```
+----------+
| AUM03002 |
+----------+
```

TEXT: File output is not supported.

TYPE: Error

REASON: This infile exit does not support write operation. Processing aborted.

ACTION: Do not use the infile exit for output.

REFERENCES: None

```
+----------+
| AUM03003 |
+----------+
```

TEXT: Opening DDName &1 ...

TYPE: Information

REASON: Record the name of the data set opened.

ACTION: None required.

REFERENCES: None

```
+----------+
| AUM03004 |
+----------+
```

TEXT: Variable &1 size is different from defined: &2.

TYPE: Warning

REASON: The SAS variable, &1, is different from what it should be, based on the header record. Its value may be truncated if shorter.

ACTION: Contact CA Technical Support.

REFERENCES: None

```
+----------+
| AUM03005 |
+----------+
```

TEXT: Mismatched Type or Internal/External: &1.

TYPE: Warning

REASON: The data type of the SAS variable was defined as numeric but the data header indicated that it should be character, or vice versa.

ACTION: Contact CA Technical Support.

REFERENCES: None

```
+----------+
| AUM03006 |
+----------+
```

TEXT: DDName &1 closed ...

TYPE: Information

REASON: Record the name of the data set.

ACTION: None required.

REFERENCES: None

```
+----------+
| AUM03007 |
+----------+
```

TEXT: Logical record length exceeds &1 bytes.

TYPE: Error

REASON: The logical record length of the data presented for processing exceeds &1 bytes, the internal buffer size for holding the data. Processing aborted.

ACTION: Verify the data export/transfer process. This may be caused by the data being shifted, or invalid record format.

REFERENCES: None

```
+----------+
| AUM03008 |
+----------+
```

TEXT: Fail to open the variable table: &1.

TYPE: Error

REASON: The variable table &1 either was not found or could not be opened for read.

ACTION: Verify that the variable table exists in the complex and there is no contention for the table.

REFERENCES: None

```
+-----------+
| AUM03009 |
+-----------+

   TEXT:  Attempt to write over an unprocessed record.

   TYPE:  Error

   REASON:  The system attempted to copy a record to the
            input buffer, but doing so would overwrite an
            unprocessed record.  Processing terminates.

   ACTION:  Contact CA Technical Support.

   REFERENCES:  None

+-----------+
| AUM03010 |
+-----------+

   TEXT:  Unsupported Data Type: &1.

   TYPE:  Warning

   REASON:  The data type &1 is not supported.

   ACTION:  None required.

   REFERENCES:  None

+-----------+
| AUM03011 |
+-----------+

   TEXT:  Scale factor for item &1 contains a negative value
            (&2).  Please report this message to CA MICS
            Product Support.

   TYPE:  Error

   REASON:  The MeasureWare data header contains a negative
            value for item &1.  Processing terminated.

   ACTION:  Contact CA Technical Support.

   REFERENCES:  None

+-----------+
| AUM03012 |
+-----------+

   TEXT:  Invalid negative value(s) in the MWA record #&1 at
            &2:&3 (row:column)

   TYPE:  Warning

   REASON:  Negative values are not normal in MeasureWare
            data.

   ACTION:  Contact HP MeasureWare Product Support

   REFERENCES:  A record dump follows this message.

+-----------+
| AUM03013 |
+-----------+

   TEXT:  Unrecognized SAS data type: &1.

   TYPE:  Error

   REASON:  The data type &1 is neither numeric nor
            character.

   ACTION:  Contact CA Technical Support.

   REFERENCES:  None

+-----------+
| AUM03014 |
+-----------+

   TEXT:  Invalid option: &1 in the INFILE statement.

   TYPE:  Error

   REASON:  While parsing the INFILE statement, an invalid
            option, &1, was found.

   ACTION:  Correct the INFILE statement.

   REFERENCES:  None

+-----------+
| AUM03015 |
+-----------+

   TEXT:  Message count for &1 exceeded.  Further messages
            will be suppressed.

   TYPE:  Information

   REASON:  The number of messages exceeds the threshold set.
```

ACTION: None required.

REFERENCES: None

```
+----------+
| AUM03016 |
+----------+
```

TEXT: &1 = &2/&3 (input from MWA record #&6, columns &4-&5)

TYPE: Error

REASON: Message 3012 noted that a negative value was detected. This message names the CA MICS element (&1), its input value (&2), and the scale factor (&3).

ACTION: Verify that you are processing data from an operating system that is compatible with one of the supported operating systems listed in Section 1.4 of this guide. If there is a mismatch between the scopeux, utility, and extract versions on your system, bring them to the same level. If the versions match those listed in Section 1.4, contact CA Technical Support.

REFERENCES: None

```
+----------+
| AUM03017 |
+----------+
```

TEXT: Work area with size &1 not allocated.

TYPE: Error

REASON: The required amount of work space was not allocated by SAS. This could be caused by system memory shortage.

ACTION: Try again later.

REFERENCES: None

```
+----------+
| AUM03018 |
+----------+
```

TEXT: The following element(s) holds SAS missing value due to its input field containing negative value:

TYPE: Warning

REASON: At least one negative value was detected during the input process. This message lists the CA MICS element whose input value is negative.

ACTION: See AUM03016

REFERENCES: None

```
+----------+
| AUM03019 |
+----------+
```

TEXT: Fail to load the variable table: &1. Metric number is out of range: &2.

TYPE: Error

REASON: An invalid metric number (&2) in the variable table &1 is detected.

ACTION: Verify that the variable table is not modified. Correct the variable table and try again.

REFERENCES: None

```
+----------+
| AUM03020 |
+----------+
```

TEXT: Input record size is bigger than buffersize: &1.

TYPE: Error

REASON: The record being read exceeds the size of the input buffer. Processing aborted.

ACTION: Verify the data transfer process. The record size for the z/OS data set must be smaller than &1.

REFERENCES: None

```
+----------+
| AUM03021 |
+----------+
```

TEXT: Record type: &1 is expected, but instead &2 is found.

```
|    TYPE:  Error                                                          |
|                                                                          |
|    REASON:  The input data's integrity is questionable.                  |
|             Processing aborted.                                          |
|                                                                          |
|    ACTION:  Verify the data export/transfer process.                     |
|                                                                          |
|    REFERENCES:  None                                                     |
|                                                                          |
| +----------+                                                             |
| | AUM03022 |                                                             |
| +----------+                                                             |
|                                                                          |
|    TEXT:  &1 contains invalid information: &2                            |
|                                                                          |
|    TYPE:  Error                                                          |
|                                                                          |
|    REASON:  Some record in the variable table &1 contains                |
|             invalid data.                                                |
|                                                                          |
|    ACTION:  Correct the record in the variable table.                    |
|                                                                          |
|    REFERENCES:  None                                                     |
|                                                                          |
| +----------+                                                             |
| | AUM03023 |                                                             |
| +----------+                                                             |
|                                                                          |
|    TEXT:  Maximum LAN devices supported (&1) exceeded: &2                |
|                                                                          |
|    TYPE:  Error                                                          |
|                                                                          |
|    REASON:  The number of LAN devices supported is &1; the               |
|             number of LAN devices noted in the data is &2.               |
|                                                                          |
|    ACTION:  Contact CA Technical Support.                                |
|                                                                          |
|    REFERENCES:  None                                                     |
|                                                                          |
| +----------+                                                             |
| | AUM03024 |                                                             |
| +----------+                                                             |
|                                                                          |
|    TEXT:  Metrics array capacity exceeded.  Data type &1 is              |
|             not supported.                                               |
|                                                                          |
|    TYPE:  Error                                                          |
|                                                                          |
|    REASON:  Verify that no user modifications were made to               |
|             your sharedprefix.MICS.SOURCE(AUMMWAV).                      |
|                                                                          |
|    ACTION:  If user modifications are found, remove them from            |
|
```

```
|             AUMMWAV. The data type is not supported by the               |
|             Analyzer.  Restart processing after removing                 |
|             the user modifications.                                      |
|                                                                          |
|             If user modifications were not found, contact                |
|             CA Technical Support.                                        |
|                                                                          |
|    REFERENCES:  None                                                     |
|                                                                          |
| +----------+                                                             |
| | AUM03025 |                                                             |
| +----------+                                                             |
|                                                                          |
|    TEXT:  Invalid record length: &1                                      |
|                                                                          |
|    TYPE:  Error                                                          |
|                                                                          |
|    REASON:  The input data contains a record with an invalid             |
|             record length &1.  Input data's integrity is in              |
|             doubt.                                                       |
|                                                                          |
|    ACTION:  Verify the data export/transfer process.                     |
|                                                                          |
|    REFERENCES:  None                                                     |
|                                                                          |
| +----------+                                                             |
| | AUM03026 |                                                             |
| +----------+                                                             |
|                                                                          |
|    TEXT:  Invalid metric number in the header record: &1                 |
|                                                                          |
|    TYPE:  Error                                                          |
|                                                                          |
|    REASON:  The header record contains an invalid metric                 |
|             number.                                                      |
|                                                                          |
|    ACTION:  Contact CA Technical Support.                                |
|                                                                          |
|    REFERENCES:  None                                                     |
|                                                                          |
| +----------+                                                             |
| | AUM03027 |                                                             |
| +----------+                                                             |
|                                                                          |
|    TEXT:  Error in Loading MWA table                                     |
|                                                                          |
|    TYPE:  Error                                                          |
|                                                                          |
|    REASON:  Errors encountered when the infile exit attempts             |
|             to load MWA table.                                           |
|                                                                          |
|    ACTION:  See prior messages for details.                              |
|
```

```
REFERENCES:  None

+----------+
| AUM03028 |
+----------+

TEXT:  Error occurs when processing input data

TYPE:  Error

REASON:  Errors encountered when the infile exit attempts
         to read the input data.

ACTION:  See prior messages for details.

REFERENCES:  None

+----------+
| AUM03029 |
+----------+

TEXT:  Total number of metrics exceeds maximum: &1
         AUM will not process the rest of metrics.

TYPE:  Error

REASON:  Header record processing uncovers that the number
         of metrics is beyond the capacity supported by
         the infile exit.

ACTION:  See prior messages for details.

REFERENCES:  None

+----------+
| AUM03099 |
+----------+

TEXT:  &1

TYPE:  Information

REASON:  Infile exit prints a line in SASLOG.

ACTION:  None required.

REFERENCES:  None
```

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

APPENDIX B

DATA DICTIONARY

```
 CCCCCCCC       AAAAAAA              MMM       MMM  IIIIIIIIII   CCCCCCCC    SSSSSSSS
CCCCCCCCCCC    AAAAAAAAA             MM MM   MM MM      III     CCCCCCCCCC  SSSSSSSSSS
CCC           AA       AA            MM  MM MM  MM      III     CCC         SS      SS
CCC           AA       AA            MM   MMM   MM      III     CCC         SS
CCC           AAAAAAAAAAA            MM    M    MM      III     CCC          SSSSSSSSS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC                 SS
CCC           AA       AA            MM         MM      III     CCC          SS     SS
CCCCCCCCCCC   AA       AA            MM         MM      III     CCCCCCCCCC   SSSSSSSSS
 CCCCCCCC     AA       AA            MM         MM  IIIIIIIIII   CCCCCCCC     SSSSSSS
```

r14.2

C A M I C S A N A L Y Z E R O P T I O N F O R M E A S U R E W A R E G U I D E

APPENDIX B

DATA DICTIONARY

Appendix B. DATA DICTIONARY

The MeasureWare Information Area is supported by the CA MICS Analyzer for MeasureWare. The following sections describe the files supported in this information area:

1 - Common Data Elements
2 - The UNIX Global Performance (AUMGPR) File
3 - The UNIX Application Performance (AUMAPP) File
4 - The UNIX Process Activity (AUMKPS) File
5 - The UNIX Process Exception (AUMIPS) File
6 - The UNIX LAN Activity (AUMNIF) File
7 - The UNIX Disk Activity (AUMDDA) File
8 - The UNIX System and Collector Profile (AUMMCF) File
9 - The UNIX Logical Volume Activity (AUMVLM) File
10 - The UNIX Business Transaction SLO Activity (AUMTTA)

B.1 - Common Data Elements

Common data elements are named for the measurement descriptive quality that they possess. These data elements may be carried as data elements on one or more of the Analyzer's files. There are also some common data elements that exist in one file (or no files), but were made common elements so that they could be given familiar names (rather than following the CA MICS convention of using the first three characters as the file prefix).

Most of these data elements are particularly interesting because they exist at the detail level of each UNIX file and are available for use by the Analyzer's user routines for account code and application name derivations.

Definitions of the Analyzer's common data elements follow, in alphabetical order.

A U M A C T x — Account Code Structure

The AUMACTx data elements are user-defined fields that define your installation's accounting or cost center structure. The first element with this alias is AUMACT1.

TYPE: Common, retained - Character, length 1 to 30.

SEE ALSO: Not applicable

ELEMENT DERIVATION: AUMACTx is set using the account code routines you code in sharedprefix.MICS.PARMS(AUMACCT), which defines the number of account codes, and sharedprefix.MICS.PARMS(AUMACRT), which defines the content of each account code.

RANGE OF VALUE: AUMACTx may contain any combination of 1 to 30 alphanumeric characters per field, according to your definition.

SUMMARIZATION PROCESS: AUMACTx is not summarized. AUMACTx is used to determine a unique key for records within a file based on the time-span, so other elements in the file can be summarized.

USAGE NOTES: There may be up to nine account codes, each representing a level of accountability for the installation. The most important account code is AUMACT1, the second most important is AUMACT2, and so forth. You define these elements in sharedprefix.MICS.PARMS(AUMACCT) and include them in the database when you install the CA MICS Analyzer for

MeasureWare OS Agents.

| A U M A P N U M | Application Number

AUMAPNUM contains the application number as defined in your Unix system.

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: AUMLDEV - Logical Device Number
AUMPIN - Process Identification Number
AUMPPID - Parent Identification Number
AUMGRPID - Group Identification Number
AUMPRI - Dispatch Priority

ELEMENT DERIVATION: AUMAPNUM is set to the value of the MeasureWare Agent Application metric APP_NUM.

RANGE OF VALUE: AUMAPNUM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: AUMAPNUM contains the last value of AUMAPNUM processed in the level of summarization.

USAGE NOTES: None

| A U M A P U | Application Name

By default, AUMAPU contains the application name as specified in the MeasureWare Agent's parm file. If changes to the default are needed, perhaps to make application names clearer in reports, you can define AUMAPU in prefix.MICS.PARMS(AUMAURT).

TYPE: Common, retained - Character, length 20.

SEE ALSO: AUMAPNUM - Application Number

ELEMENT DERIVATION: AUMAPU is set to the value of the MeasureWare Agent Application metric APP_NAME.

RANGE OF VALUE: AUMAPU can contain any combination of 1 to 20 alphanumeric characters. restrictions:

SUMMARIZATION PROCESS: AUMAPU is not summarized. AUMAPU is used to determine a unique key for records within a file based on the time-span, so other elements in the file can be summarized.

USAGE NOTES: None

| A U M C P M D L | System Model Identification

AUMCPMDL contains the model number of the Unix system from which the data was collected.

TYPE: Common, retained - Character, length 12.

SEE ALSO: AUMOPSYS - Operating System Name
AUMOPSVR - Operating System Release Number
AUMOPSLC - Operating System License

ELEMENT DERIVATION: AUMCPMDL is set to the model number of the Unix system from which the data was collected.

RANGE OF VALUE: AUMCPMDL can contain any combination of 1 to 12 alphanumeric characters.

SUMMARIZATION PROCESS: AUMCPMDL contains the last value of AUMCPMDL processed in the level of summarization.

USAGE NOTES: None

| A U M G R P I D | Group Identification Number

AUMGRPID contains the group identification number taken from the MeasureWare Agent password file when the process was created.

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: AUMAPNUM - Application Number
AUMLDEV - Logical Device Number
AUMPIN - Process Identification Number
AUMPPID - Parent Identification Number
AUMPRI - Dispatch Priority

ELEMENT DERIVATION: AUMGRPID is set to set to the value of the MeasureWare Agent Process metric PROC_GROUP_ID.

RANGE OF VALUE: AUMGRPID takes any numeric value ranging from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: AUMGRPID contains the last value of AUMGRPID processed in the level of summarization.

USAGE NOTES: AUMGRPID represents the effective group identification number of the process on the HP-UX and AIX MeasureWare platforms and the real group identification number of the process on the SOL platform.

A U M L D E V | Logical Device Number

AUMLDEV contains the logical device number for a disk drive.

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: MCFNDSKS - Number of Disk Drives Configured

ELEMENT DERIVATION: AUMLDEV is set to the value of the MeasureWare Agent Disk metric BYDSK_ID.

RANGE OF VALUE: AUMLDEV takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: AUMLDEV contains the last value of AUMLDEV processed in the level of summarization.

USAGE NOTES: None

A U M L G V E R | Measurement Interface Version Number

AUMLGVER contains the log file's version number.

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: AUMSCOPE - Scope Collector Name and Version Number

ELEMENT DERIVATION: AUMLGVER is set to the value of the MeasureWare Agent Configuration metric GBL_LOGFILE_VERSION.

RANGE OF VALUE: AUMLGVER takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: AUMLGVER contains the last value of AUMLGVER processed in the level of summarization.

USAGE NOTES: None

A U M L N T T Y | Logon TTY Name

AUMLNTTY contains the device file name of the login or spawning port.

TYPE: Common, retained - Character, length 4.

SEE ALSO: Not applicable

ELEMENT DERIVATION: AUMLNTTY is set to the value of the MeasureWare Agent Process metric PROC_TTY.

RANGE OF VALUE: AUMLNTTY can contain any combination of 1 to 4 alphanumeric characters.

SUMMARIZATION PROCESS: AUMLNTTY is not summarized. AUMLNTTY is used to determine a unique key for records within a file based on the timespan, so other elements in the file can be summarized.

USAGE NOTES: AUMLNTTY is blank if there is no controlling terminal for the process.

A U M O P S L C | Operating System License

AUMOPSLC contains the character code indicating the number of users authorized to use this operating system license.

TYPE: Common, retained - Character, length 12.

SEE ALSO: AUMOPSSVR- Operating System Release Number
AUMOPSYS - Operating System Name
AUMCPMDL - System Model Identification

ELEMENT DERIVATION: AUMOPSLC is set to the value of the MeasureWare Agent Configuration metric GBL_OSVERSION.

RANGE OF VALUE: AUMOPSLC can contain any combination of 1 to 12 alphanumeric characters.

SUMMARIZATION PROCESS: AUMOPSLC contains the last value of AUMOPSLC processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| A U M O P S V R |  Operating System Release Number
+-----------------+
AUMOPSVR contains the operating system release number.

TYPE:  Common, retained - Character, length 12.

SEE ALSO:  AUMOPSYS - Operating System Name
           AUMOPSLC - Operating System License
           AUMCPMDL - System Model Identification

ELEMENT DERIVATION:  AUMOPSVR is set to the value of the
MeasureWare Agent Configuration metric GBL_OSRELEASE.

RANGE OF VALUE:  AUMOPSVR can contain any combination of 1 to
12 alphanumeric characters.

SUMMARIZATION PROCESS:  AUMOPSVR contains the last value of
AUMOPSVR processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A U M O P S Y S |  Operating System Name
+-----------------+
AUMOPSYS contains the operating system name (HP-UX, for
example).

TYPE:  Common, retained - Character, length 12.

SEE ALSO:  AUMOPSVR - Operating System Release Number
           AUMOPSLC - Operating System License
           AUMCPMDL - System Model Identification
           MCFOPSYS - Operating System Name

ELEMENT DERIVATION:  AUMOPSYS is set to the value of the
MeasureWare Agent Configuration metric GBL_OSNAME.

RANGE OF VALUE:  AUMOPSYS can contain any combination of 1 to
12 alphanumeric characters.

SUMMARIZATION PROCESS:  AUMOPSYS is not summarized. AUMOPSYS
is used to determine a unique key for records within a file
based on the time-span, so other elements in the file can be
summarized.

USAGE NOTES:  None
```

```
+-----------------+
| A U M P I N     |  Process Identification Number
+-----------------+
AUMPIN contains the process identification number (PID, in
Unix terms).

TYPE:  Common, retained - Numeric, measurement unit is
integer.

SEE ALSO:  AUMAPNUM - Application Number
           AUMLDEV  - Logical Device Number
           AUMPPID  - Parent Identification Number
           AUMGRPID - Group Identification Number
           AUMPRI   - Dispatch Priority

ELEMENT DERIVATION:  AUMPIN is set to the value of the
MeasureWare Agent Process metric PROC_PROC_ID.

RANGE OF VALUE:  AUMPIN takes any numeric value ranging from
0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  AUMPIN is not summarized. AUMPIN is
used to determine a unique key for records within a file
based on the time-span, so other elements in the file can be
summarized.

USAGE NOTES:  None


+-----------------+
| A U M P P I D   |  Parent Identification Number
+-----------------+
AUMPPID contains the parent's process identification number
(PID).

TYPE:  Common, retained - Numeric, measurement unit is
integer.

SEE ALSO:  AUMAPNUM - Application Number
           AUMLDEV  - Logical Device Number
           AUMPIN   - Process Identification Number
           AUMGRPID - Group Identification Number
           AUMPRI   - Dispatch Priority

ELEMENT DERIVATION:  AUMPPID is set to the value of the
MeasureWare Agent Process metric PROC_PARENT_PROC_ID.

RANGE OF VALUE:  AUMPPID takes any numeric value ranging from
0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  AUMPPID contains the last value of
AUMPPID processed in the level of summarization.
```

USAGE NOTES: None

| A U M P R I | Dispatch Priority

AUMPRI contains the process's dispatch priority.

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: AUMAPNUM - Application Number
AUMLDEV - Logical Device Number
AUMPPID - Parent Identification Number
AUMGRPID - Group Identification Number
AUMPIN - Process Identification Number

ELEMENT DERIVATION: AUMPRI is set to the value of the MeasureWare Agent Process metric PROC_PRI for HP-UX, or PROC_REVERSE_PRI for Solaris.

RANGE OF VALUE: AUMPRI takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: AUMPRI contains the last value of AUMPRI processed in the level of summarization.

USAGE NOTES: The lower the dispatch priority, the more likely the process will be dispatched. In HP-UX systems, a real time process has a dispatch priority between 0 and 127. Values above 127 are normal priorities that the kernel can change to balance the load.

| A U M P R O G | Program Name

AUMPROG contains the process program name.

TYPE: Common, retained - Character, length 16.

SEE ALSO: AUMUSER - User Logon Name

ELEMENT DERIVATION: AUMPROG is set to the value of the MeasureWare Agent Process metric PROC_PROC_NAME.

RANGE OF VALUE: AUMPROG can contain any combination of 1 to 16 alphanumeric characters.

SUMMARIZATION PROCESS: AUMPROG is not summarized. AUMPROG is used to determine a unique key for records within a file based on the time-span, so other elements in the file can be summarized.

USAGE NOTES: None

| A U M S A M P L | Number of Data Samples

AUMSAMPL contains the number of samples taken during the measurement interval.

TYPE: Common, accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: AUMSAMPL is set to the value of the MeasureWare Agent Global metric GBL_PROC_SAMPLE.

RANGE OF VALUE: AUMSAMPL takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: AUMSAMPL contains the sum of all values of AUMSAMPL processed in the level of summarization.

USAGE NOTES: None

| A U M S C O P E | Scope Collector Name and Version Number

AUMSCOPE contains the Scope collector name and version number. For example, SCOPE/UX B.02.00(41), where "B" is the version, "02" is the update level, "00" is the fix level, and "41" is HP's fix designator.

TYPE: Common, retained - Character, length 24.

SEE ALSO: AUMLGVER - Measurement Interface Version Number

ELEMENT DERIVATION: AUMSCOPE is set to the value of the MeasureWare Agent Configuration metric GBL_COLLECTOR.

RANGE OF VALUE: AUMSCOPE can contain any combination of 1 to 24 alphanumeric characters.

SUMMARIZATION PROCESS: AUMSCOPE contains the last value of AUMSCOPE processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| A U M S Y S I D |  Name of Host System
+-----------------+
AUMSYSID contains the Unix system's name.

TYPE:  Common, retained - Character, length 40.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  AUMSYSID is set to the value of the
MeasureWare Agent Configuration metric GBL_SYSTEM_ID.

RANGE OF VALUE:  AUMSYSID can contain any combination of 1 to
40 alphanumeric characters.

SUMMARIZATION PROCESS:  AUMSYSID contains the last value of
AUMSYSID processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A U M U S E R   |  User Logon Name
+-----------------+
AUMUSER contains the user's login name.

TYPE:  Common, retained - Character, length 16.

SEE ALSO:  AUMPROG  - Program Name

ELEMENT DERIVATION:  AUMUSER is set to the value of the
MeasureWare Agent Process metric PROC_USER_NAME.

RANGE OF VALUE:  AUMUSER can contain any combination of 1 to
16 alphanumeric characters.

SUMMARIZATION PROCESS:  AUMUSER is not summarized. AUMUSER
is used to determine a unique key for records within a file
based on the time-span, so other elements in the file can be
summarized.

USAGE NOTES:  None


+-----------------+
| D A Y           |   Day of Month
+-----------------+
The calendar day of the month, a number from 1 to 31.

TYPE: Common, retained - Numeric, measurement unit is
integer.


SEE ALSO:  ENDTS - End Time Stamp

ELEMENT DERIVATION: DAY is derived from the standard CA MICS
time-stamp field ENDTS.

If the 13MONTHYEAR option of MICS.PARMS(SITE) is used to
define your fiscal year as having 13 months of 28 days each,
then the value of DAY is set according to how many days from
the beginning of one of the thirteen "months" the day
represented by ENDTS falls.

For example, if your fiscal year is specified as starting on
July 1, then July 28 is DAY=28 of MONTH=01, while July 29 is
DAY=01 of MONTH=02.  The thirteenth "month" will have 29 days
in a non-leap year and 30 days in a leap year.

RANGE OF VALUE: DAY takes any numeric value in the following
ranges:

  1 to 31

  1 to 30 if the 13MONTHYEAR  option is in use

SUMMARIZATION PROCESS:  DAY contains the value of the last
DAY encountered.  It is not meaningful above the DAYS
timespan.

USAGE NOTES:  DAY can be used to select data for a day in
which you are interested.  Remember that files in the DETAIL
and DAYS timespans may contain data for more than one day,
and therefore selection is needed.


+-----------------+
| D A Y N A M E   |   Name of Day of Week
+-----------------+
DAYNAME is the name of the day of the week associated with
the measurement data in a CA MICS record.

TYPE: Common, retained, - Character, length 3.

SEE ALSO:  DAY   - Day of Month
           ENDTS - End Time Stamp

ELEMENT DERIVATION: DAYNAME is derived from the standard
CA MICS time-stamp field ENDTS.

RANGE OF VALUE:  DAYNAME can contain any of the following:

  SUN, MON, TUE, WED, THU, FRI, or SAT.

SUMMARIZATION PROCESS:  DAYNAME contains the value of the
```

last DAYNAME encountered. DAYNAME is not meaningful above the DAYS level.

USAGE NOTES: You can use DAYNAME to select data applicable to a day of the week in which you are interested. Remember that a cycle in the DETAIL or DAYS timespan may contain data from more than one day, so selection is necessary.

D U R A T I O N | Recording Interval Time

DURATION contains the number of seconds in the "reporting interval" of a given monitor. Reporting interval is the period of time during which a monitor gathers and reports statistics.

DURATION is calculated for the monitors RMF, TSO/MON, CICS, IMS, NPM, NetView, CA NetSpy, CA NetMaster, DB2, VM, VCC/VCA, TAP, and MQSeries. In these monitors, DURATION is the length of time that one record covers in the DETAIL timespan. In higher timespans, DURATION is the length of time during which the data collector provides measurement data within the interval that is covered by the CA MICS record for the time period.

The length of the RMF reporting interval is specified via the RMF INTERVAL parameter. Allowed values are from 1 to 60 minutes. The default value for the RMF data recording interval is 30 minutes for the RMF record types 70-79. The Hardware Instrumentation Services (HIS) create SMF type 113 records, also processed by the CA MICS Hardware and SCP Analyzer Option. The duration for SMF type 113 interval records is 15 minutes.

The length of the TSO/MON reporting interval is specified via the INTVL parameter of the OPTIONS macro. Five minutes (INTVL=5) is the recommended value.

The length of a CICS recording interval is specified for CMF (CICS Monitoring Facility) at CICS/MVS Version 2 via the FREQ parameter in DFHMCT TYPE=RECORD macro. This affects the duration of the performance class global records. For CMF input at CICS/ESA Release 3.1.1 and higher, DURATION is set to the recording interval of CICS statistics data. You can change the statistics recording interval in CICS using the CEMT SET STATISTICS command. For TMON/CICS, the system interval records are written on a minute basis.

For IMS, you specify the number of IMS log records to record between IMS checkpoints. After these records are written to the IMS log, an IMS checkpoint is taken. Recognition of this checkpoint is the event CA MICS uses to produce another DETAIL IMS system record. You specify the number of log records between checkpoints with the CPLOG= parameter of the IMSCTF macro in the IMS Stage I generation. Note that under this scheme for logging checkpoints to the IMS log, DURATION is not a constant value for all CA MICS IMS DETAIL observations, and may, in fact, vary greatly according to IMS system load and time of day.

NPM, the NetView Performance Monitor, is divided into two subsystems: the Network Subsystem and the Session Subsystem. For the Network Subsystem, the length of the NPM recording interval is specified via the INTERVAL operand of the NPM INTERVAL initialization parameter and the online START COLLECT command. For the NPM Session Subsystem, the length of the interval is specified via the SESSINT operand of the NPM INTERVAL initialization parameter.

For NetView Session Monitor/NLDM, interval recording is provided through a Network Communications Control Facility (NCCF) CLIST. The CLIST issues an NCCF EVERY command to collect NLDM/RTM data on the interval you choose.

TCP information area files:

For all the files in the TCP area, DURATION is the elapse time between a connection start time and end time.

- For TCP Termination (subtype 2) records, DURATION is calculated by subtracting connection start time from connection end time.
- For TN3270E Telnet server SNA session termination (subtype 21) records, DURATION comes from field SMF119TN_NTDur at offset x'64' in SMF type 119.
- For TSO Telnet client termination (subtype 23) records, DURATION comes from field SMF119TN_CTDur at offset x'4C' in SMF type 119.
- For FTP client transfer completion (subtype 3) records, DURATION comes from field SMF119FT_FCDur at offset x'74' in SMF type 119.
- For FTP server transfer completion (subtype 70) records, DURATION comes from field SMF119FT_FSDur at offset x'70' in SMF type 119.
- For TCPIP Statistics (subtype 5) records that generate TCPIPX, TCPTCI, TCPUDI, and TCPICX files, the DURATION in these files comes from the respective fields SMF119AP_TSIPDuration, SMF119AP_TSTCDuration,

SMF119AP_TSUDDuration, and SMF119AP_TSICDuration at offset x'00' in SMF type 119.

- For Interface statistics (subtype 6) records, DURATION comes from field SMF119IS_IFDuration at offset x'00' in SMF type 119.

CDZ information area files:

For the CCT IBM Sterling Connect:Direct file, DURATION is th elapse time between time the COPY step started (STARTS) and ENDTS.

In DB2, the DURATION element for the DB2DSY and DB2DSD files represents the length of time that is reported in each detail observation. DB2 Release 2 provided DB2 systems personnel with the ability to define, at DB2 system startup, the time interval in which SMF system statistics are produced. This value is specified in minutes and should equal DURATION.

For VCA, except for the exceptions described below, DURATION contains the difference in time between the previous and current executions of the CA MICS Space Collector Option (VCC) scanning process.

For the first execution of a VCC scan, no previous scan timestamp is available. In this case, DURATION is set to the value on the OPTIONS statement specified in prefix.MICS.PARMS(VCAOPS), default 24 hours.

For data sets that do not have a data set creation time, whose create date is later that the previous VCC scan date, DURATION is calculated by multiplying the number of days the data set existed prior to the current VCC scan, plus the time of day when the current VCC scan executed.

Beginning with z/OS 1.11, the Volume Table of Contents (VTOC) for Extended Address Volumes (EAV) began recording both the date and time of data set creation. For data sets with both the time and date of creation, if they were created after the previous VCC scan, DURATION is calculated as the difference between the current VCC scan time and data set creation time.

For data sets that are created after the VCC scan started, the DURATION value is set to zero.

For VM/370 Monitor, the recording interval length is specified via the INTERVAL operand of the MONITOR command. For VM/XA or VM/ESA Monitor, the interval length is specified via the INTERVAL operand of the MONITOR SAMPLE command.

In MQSeries, the DURATION element for the MQSBMS, MQSDMS, MQSLMS, MQSMMS, MQSCMS, and MQSDBS files represents the length of time that is reported in each detail observation.

The Message Queue Manager (MQM) allows an interval length specification at MQM system startup via the STATIME parameter. If not specified, the system default is 30 minutes. In the MQ Message Broker MQBMFA and MQBMFN files, the DURATION represents the difference between the ENDTS and STARTS of the message flow interval.

In all files for the CA MICS Analyzer Option for CA IDMS, except the IDMS System Activity File (IDMSSY), DURATION represents the elapsed time of the individual transaction records. In summarized timespans, DURATION represents the sum of DURATION of the individual transaction records that contributed to the summarized observation. DURATION is calculated as the difference between the End Date and Time (TASEDATE and TASETIME) and the Start Date and Time (TASSDATE and TASSTIME) fields from the CA IDMS Performance Monitor header record.

In all files for the CA MICS Analyzer Option for CA IDMS except the IDMS System Activity File (IDMSSY), the Start Time Stamp (STARTTS), and End Time Stamp (ENDTS), data elements reflect the start and end times of the individual transactions. For IDMSSY file, STARTTS and ENDTS are adjusted to represent the start and end of the time interval within which the transactions executed.

TYPE: Common accumulated data element - NUMERIC, measurement unit is INTEGER.

SEE ALSO: ENDTS - End Timestamp
STARTTS - Start Timestamp
INTERVLS - Number of Recording Intervals
SAMPLES - Number of Samples

ELEMENT DERIVATION: For RMF record types 70-79, DURATION contains the value input from field SMFxxINT. It is computed according to the following formula:

```
DURATION=60*FLOOR(SMFxxINT/100000)+MOD(SMFxxINT,100000)
/1000;
```

where SMFxxINT is in the form "MMSSTTT"

For the RMF type 113 record, DURATION is computed by subtracting STARTTS from ENDTS, input from the SMF type 113 record fields SMF113IntervalStart and SMF113IntervalEnd, respectively.

For TSO/MON, DURATION contains the value input from field

TSMSINTV of the System Record (usually assigned SMF number 200).

For CICS, DURATION is measured from interval start to end based on timestamps in the record: DURATION=ENDTS-STARTTS. However, for CMF processing at CICS/ESA Release 3.1.1 and higher, DURATION is set to the sum of dispatch time (CSYQRDSP) and wait time (CSYQRWTM) when the statistics record is written for either the first or last (EOD) recording interval. The derivation is: DURATION=SUM(CSYQRDSP,CSYQRWTM).

For IMS, DURATION is calculated as described below unless the system interval time specification (from the IMSOPS member of the prefix.MICS.PARMS library on the OPTIONS statement) has been exceeded by 50 percent and the IMS checkpoint number of the previous checkpoint is not equal to the current checkpoint number-1. If these conditions are met, DURATION is set to the IMSOPS interval specification.

Under normal circumstances, the IMS DURATION contains the difference between the end timestamp of the current checkpoint and the end timestamp of the previous checkpoint. The interval specification on the IMSOPS OPTIONS statement does not override this calculation if the IMS checkpoint numbers monotonically increase (for example, 1, 2, 3, ...).

For example, assume that the specified system interval is 15 minutes. If the initial calculation of DURATION was 17 minutes, DURATION would contain a value of 17 minutes. If the initial calculation was 22.6 minutes (greater than 1.5 times 15 minutes, or 22.5 minutes), DURATION would be set to 15 minutes, except for IMS systems with monotonically increasing checkpoint numbers.

For the NPM Network Subsystem, DURATION contains the value input from field LOGTIME, which is the actual NPA observation time for the NCP being used. For the NPM Session Subsystem, DURATION contains the difference between the interval end timestamp, VPVEINT, and the interval start timestamp, VPVDT.

For NLDM, DURATION contains the difference between the Collection Period Begin Time Stamp, LRTMCOLB, and the Collection Period End Time Stamp, LRTMCOLE.

For CA NetMaster, DURATION equals the elapsed time between session start and session end time.

For DB2 interval records, DURATION contains the difference between the end timestamp of the current SMF record and the end timestamp of the previously processed SMF record.

For VCC and VCA, DURATION is the difference between the time of the previous VCC scan execution, which is saved in the VCC control file, and the time of the current VCC scan execution. Exceptions to this computation, described earlier, occur when processing records with an initial VCC scan, and newly created data sets.

For the TLM component of the CA MICS Tape Analyzer Option, DURATION is the difference between the time of the previous TMCC scan execution, which is saved in the TMCC control file, and the time of the current TMCC scan execution.

For the VTS component of the CA MICS Tape Analyzer Option:

For the following SVT StorageTek VSM Virtual Tape Subsystem Information Area (SVT) DETAIL timespan files, DURATION is set to 900 seconds (15 minutes):

- o VTSS Subsystem Performance (SVTSSP) file
- o Channel Interface Performance (SVTCIP) file
- o Real Tape Drive Performance (SVTRPR) file
- o StorageTek VTSS Summary (SVTSVS) file

For the IBM Virtual Tape Subsystem Information (IVT) Area DETAIL and DAYS timespan files, DURATION is set to exactly 3600 seconds (1 hour).

For the IBM Hydra Virtual Tape Server Information (HVT) Area, DURATION contains the value input from field "Interval Duration" in the header section of each input record.

For the VM component, VMCCMU obtains DURATION from type 01 and type C1 account records columns 29-32. VMCDAC gets DURATION from types 02, C2, 03, and C3 account records columns 29-32. In VMCVMU file, DURATION contains the difference between VMUENDTS and STARTTS. In other VMC and VMX Information Area files, DURATION contains the difference between ENDTS and STARTTS in the DETAIL timespan.

For the CA MICS Analyzer Option for MeasureWare, DURATION is set to the value of the MeasureWare Global metric INTERVAL.

For the CA MICS Analyzer Option for MQSeries, DURATION contains the difference between the end timestamp of the current SMF record and the end timestamp of the previously processed SMF record. In the MQ Message Broker MQBMFA and MQBMFN files, the DURATION represents the difference between the ENDTS and STARTS of the message flow interval - the interval end datetime and interval start datetime.

For the CA MICS Analyzer Option for CA IDMS:

For the CA IDMS files, DURATION contains the difference between the ENDTS and STARTTS data elements. STARTTS and ENDTS contain the values of the start date and time (TASSDATE and TASSTIME) and end date and time (TASEDATE and TASETIME) fields from the CA IDMS Performance Monitor header record.

In individual IDMSUA file transaction records, DURATION is recorded in one of six category response time data elements (conversational, long, medium, short, excessive, and user-defined), based on the transaction type. Transaction type is determined in user written code in the relative longevity exit IDMRLRT.

RANGE OF VALUE: The value of this element ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DURATION contains the sum of all occurrences of DURATION that were included in the level of summarization.

USAGE NOTES: At the DETAIL level, DURATION is no larger than ENDTS-STARTTS, but may be smaller if the data collector being used does not provide measurement data for all of the time between STARTTS and ENDTS.

For VCA, DURATION contains the time interval between VCC scan executions. However, if VCC processing is performed frequently, or if two or more devices are processed, the computed values may not provide the results you expect for computed average values on reports. It is difficult to report average space utilization on volumes if the volumes are not processed with each VCC scan execution.

E N D T S	End Timestamp

ENDTS has different definitions, depending on the timespan and the data collector.

DETAIL Timespan:

- o ENDTS is the time of termination of an executable function, such as job, command, and so on.
- o For incident files in which ENDTS exists, ENDTS is the time at which the incident occurred:

  For the CICS, IMS, VM/CMS, and Network Analyzer Option incident files, this is the end timestamp of the transaction or system measurement interval during which the incident occurred.

- o For other files in which ENDTS exists, it is the time that the data collector's measurement interval ended:

  For CA IDMS data:

  For the IDMS System Activity File (IDMSSY), ENDTS is adjusted based on a user-defined interval value. This value is coded in the INTERVAL statement of prefix.MICS.PARMS(IDMOPS) and determines the cut-off point for each timespan.

  For all other CA IDMS files, ENDTS is the date and time when the IDMS transaction completed. ENDTS contains the value of the end date and time fields (TASEDATE and TASETIME) from the #PMTAS section of the CA IDMS Performance Monitor record.

  For RMF-derived data:

  For files generated from the SMF type 70-79 RMF records, ENDTS is computed by adding the duration of the interval to the starting time for compatibility with RMF.

  For files generated from the SMF type 113 subtype 1 or 2 CPU Hardware capacity, reporting, and statistics record generated by the IBM CPU Measurement Facility (CPU MF) Hardware Instrumentation Services (HIS).

  ENDTS is set equal to field SMF113_2_CTM (subtype 2) or field SMF113IntervalEnd (subtype 1), adjusted to local time by adding the GMT Offset value.

  For TSO/MON data:

  For the TSOTSO file, ENDTS is the time (date and time) when the TSO/MON record writer routine processed the record.

  For the TSOTSI and TSOTSA files, ENDTS is the time when the TSO command or subcommand ended.

  For the TSOTSU file, CA MICS combines user information from multiple TSO/MON system records when the user was inactive for one or more recording intervals. In this case, ENDTS is the timestamp from the last record consolidated into the DETAIL-level record.

For the TSOTSC file, data is summarized by HOUR at the DETAIL level. (This is unusual for CA MICS, but is required by the volume of data.) ENDTS is taken from the timestamp of the last TSO/MON system record included in the hour.

For VCC/VCA data:

ENDTS is the time from the last execution of the VCC scan. If multiple VCC executions are processed into one update execution, then the value is the last time within that particular level of summarization.

For the TLM component of the CA MICS Tape Analyzer Option:

ENDTS is the time from the last execution of the TMCC scan. If multiple TMCC executions are processed into one update execution, then the value is the last time within that particular level of summarization.

For the VTS component of the CA MICS Tape Analyzer Option:

For the IBM Virtual Tape Subsystem Information Area (IVT) files, ENDTS is derived as follows:

First, the date and time the SMF type 94 record was written to the SMF buffer is determined. The raw metrics in the record relate to the previous hour, so the code logic establishes ENDTS at 59 minutes and 59.99 seconds of that previous hour.

For example, if an SMF type 94 record was written at 13:12:48.02, the ENDTS for all IVTfff files generated from that record would be 12:59:59.99.

For the Sun StorageTek Information Area (SVT) files, ENDTS is derived as follows:

ENDTS is the time when the Sun StorageTek user SMF record was created. It is read from bytes 3 to 10 on the SMF record.

For VM Monitor interval recording data:

ENDTS in the CA MICS Analyzer Option for VM/CMS is the time when VM Monitor created the record.

For HP MeasureWare for HP-UX Systems data:

ENDTS in the CA MICS Analyzer Option for MeasureWare is the time that the data was captured in the log file.

For WebSphere MQ data:

In all files except the MQBMFA file, ENDTS is set to the value in the SMF header that tells the time and date the record was written to the SMF buffer, for example SM116DTE and SM116TME.

In the MQBMFA file, ENDTS is set to the interval end date and time values found in the message flow accounting record fields IMFLENDT and IMFLENTM.

For Web Analyzer:

ENDTS is set to the value input from field SM1209CQ of the SMF type 120 record subtype 9 at offset 68 (x'44') of the z/OS Request Information section.

All Other Timespans:

ENDTS is set to the maximum-end timestamp encountered from the records used to compile the summarized record. Thus, the combination of STARTTS and ENDTS in the observation identify the boundaries of the range of time over which the data in the record was summarized.

In the CA MICS Accounting and Chargeback Option journal files, ENDTS corresponds to the ENDTS in the component record being charged. In the ledger files, it is summarized from the journal file records.

TYPE: Common, retained - Numeric, measurement unit is SAS datetime value.

SEE ALSO: STARTTS - Start Timestamp
DURATION - Recording Interval Time
INTERVLS - Number of Recording Intervals

ELEMENT DERIVATION:

For RMF records, CA MICS follows the RMF convention and calculates ENDTS from the "beginning of interval" datetime given in fields SMFnnDAT and SMFnnIST, and the length of the interval, SMFnnINT.

For non-RMF records, ENDTS contains the value input from fields SMFnnDTE and SMFnnTME, where "nn" is the number of the SMF record. If the SMF record number is greater than 99, the input fields are named SMnnnDTE and SMnnnTME.

In the DAYS level of the ADMEXC file, ENDTS is set to the ENDTS of the CA MICS record from which the exception condition was detected.

For the CICS files, ENDTS is set to the following:

CICS CMF: The source of this data element is raw CICS monitor input data.

The CMF timestamp fields are presented in "TOD Clock Format," and are adjusted by the GMT offset for the system being processed. This offset is specified by the user in the CICOPS member of the MICS.PARMS library, on the CICOPTS statement. The field type is T, indicating a timestamp, an 8 byte copy of the output of a local store clock (STCK) instruction.

For transaction-level records written to the User Activity file (CICCSU) Application Unit Activity file (CICCAU) and the User Application Count file (CICCAC), ENDTS is either the time at which the user-task was detached or the time at which data recording was completed in support of CONV=YES. The source of this data element for these files follows:

Monitor	Input Field	Description
CTS 2.2 to 3.1	MNR_ID_STOP	Transaction stop time
CTS 3.2 and above	STOP (FID 006)	Finish time of measurement interval in the DFHCICS group

For system-level records written to the System Activity (CICCSY) file and the Dictionary file (CICCDC), ENDTS is the time at which the recording interval ended. The source of this data element for these files follows:

Monitor	Input Field	Description
CTS 2.2 and above	SMFSTDAT and SMFSTCLT	Collection date Collection time

For records written to the Incident file (CICCIN), ENDTS is the end time of the exception's occurrence. The source of this data element for this file follows:

Monitor	Input Field	Description
CTS 2.2 and above	EXCMNSTO	Exception stop time

ASG-TMON for CICS TS (TCE): The source of this data element is raw Monitor input data.

For transaction-level records written to the User Activity file (CICCSU) Application Unit Activity file (CICCAU), the User Application Count file (CICCAC), the File/DBD Activity file (CICCSF), and the CICS MRO Activity file (CICCMR), ENDTS is the time of either the transaction's termination or the conversational task's termination. The source of this data element for these files follows:

Monitor	Input Field	Description
TCE	TAENDDTS	Transaction end date

For system-level records written to the System Activity (CICCSY) file, ENDTS is the time at which the recording interval ended. The source of this data element for this file follows:

Monitor	Input Field	Description
TCE	TIENDDTS	Interval end date

For records written to the Incident file (CICCIN), ENDTS is the last available end timestamp from the Transaction (TA) records.

For The Monitor releases 7.1 and below, ENDTS is derived from the start timestamp of the next system record since the end timestamp is not in the raw data record.

CICS Transaction Gateway: For the CTG information area files created from the SMF type 111 CICS Transaction Gateway records, ENDTS is derived from the date and time values when the raw data record was written to the SMF buffer. The source of this data element follows:

Monitor	Input Field	Description
CTG	SMF111_SMFSTDTE SMF111_SMFTME	Date and time when record was moved to SMF buffer

For the DB2 files, ENDTS is set to the following:

ENDTS is derived from the date-time value found in the SMF

header segment of the interval or thread event generated SMF record used to create the CA MICS database record.

For database records created from the interval based statistics records (SMF type 100, subtype 0 and 1), ENDTS correlates with the startup of the DB2 region and the user specified interval duration value (usually 30 minutes, but less if the region is stopped).

For database transaction-level records created from the SMF type 101, subtype 0 and 1 records, ENDTS correlates to an event, such as a thread termination, an authorization ID change in a thread re-use environment, or reaching the threshold value associated with DB2 ACCUMAC SMF record summarization.

For the IMS files, ENDTS is set to the following:

IMSISY (IMS System Activity) File: ENDTS contains the values input from the date and time fields found in the configuration checkpoint end record (type '40'X, subtype '98'X) on the IMS log.

IMSINC (IMS Incident) File: ENDTS contains the value input from the last ENDTS constructed for any other file.

Transaction-oriented Files (IMSISU User Activity, IMSIAU Application Unit Activity, and IMSIAC User Application Count Files): ENDTS depends on the collection method. With IMF, ENDTS is constructed from the end date and time presented by IMF in its DETAIL transaction accounting record (IMS log record type 'FA'X).

Without IMF, ENDTS is constructed from the date and time fields of the IMS log record that signifies the end of processing region occupancy for the message being analyzed. These fields may come from the type '31'X log record, field QLGUDTME, or the type '07'X log record, fields DLRDTE and DLRTME.

Except for those CA MICS Network Analyzer Option files listed below, ENDTS contains the value input from the SMF record end timestamp from the CA MICS Network Analyzer Option data sources. The exceptions, by file, are as follows:

- o NVSRTM and SNTNSV files - ENDTS is taken from the field LRTMCOLE of the RTM data section of the SMF type 39 record written by NetView/NLDM or CA NetMaster.

- o NVSNSA file - ENDTS is taken from field LACCENDT of the Accounting and Availability data section of the SMF type 39 record written by NetView/NLDM or CA NetMaster.

- o NVSNSC and NVSRTE files - There is only one CA MICS observation for each unique SNA session and the ENDTS for that observation is taken from the highest timestamp found in all SMF type 39 records (written by NetView/NLDM or CA NetMaster) for that session.

- o CDZCCT file - ENDTS is derived from the date-time value found in the SMF header segment of the IBM Sterling Connect:Direct copy termination record, SMF 133 (user defined), subtype CT.

For all VSE/Power files, the ENDTS is derived from the ACDATE and ACSTOP fields of the account record used to create the observation.

For the VMC and VMX files, the derivation of ENDTS depends on various sources:

- o VM and VM account data - VMCCMU, VMCDAC, VMCNDT and VMCVRA files obtain ENDTS from record types 01, C1, 02, C2, 03, C3, 04, 05, 06, 08, C0, 0C (RSCS) records columns 23-28. VMCNTU obtains ENDTS from record type 07 columns 53-60.

- o VM/ESA and z/VM monitor data - For files that have one observation per input record, the value of ENDTS is taken directly from the time-of-day clock stamp (MRHDRTOD) in the record header of input record. For other files, the value of ENDTS is taken from the record with the highest timestamp that contributes to the observation.

For CA MICS Analyzer Option for MeasureWare data, ENDTS is set to the value of the configuration metric DATE_SECONDS.

RANGE OF VALUE: ENDTS contains any valid SAS datetime value. The value of this data element is not affected by use of the 13MONTHYEAR option in prefix.MICS.PARMS(SITE).

SUMMARIZATION PROCESS: ENDTS contains the value of the last occurrence of ENDTS in the observations processed for the summarization level.

USAGE NOTES: The value of DURATION, which is the amount of time in the measurement interval, is sometimes smaller than ENDTS minus STARTTS. This occurs when measurement data was not available from a data collector for some of the time in the interval. A common mistake is to compute averages using ENDTS minus STARTTS.

```
+-----------------+
| H O U R         |   HOUR OF DAY
+-----------------+
The hour of the day, based on a 24 hour clock.

TYPE:  Common, retained - Numeric, measurement unit is
integer.

SEE ALSO:  ENDTS - End Time Stamp

ELEMENT DERIVATION:  HOUR is derived from the standard
CA MICS time-stamp field ENDTS.

RANGE OF VALUE:  HOUR takes any numeric value in the
following range:

  00 through 23, where 00 represents the hour 00:00:00.00
  through 00:59:59.99 and 23 represents the hour 23:00.00.00
  through 23:59:59.99.

SUMMARIZATION PROCESS:  HOUR contains the value of the last
HOUR encountered.  HOUR is not meaningful above the WEEKS
timespan.

At the DAYS and WEEKS levels of summarization in most CA MICS
files, HOUR is used to determine a unique key for records
within a file based on the timespan so other elements in
the file can be summarized.

USAGE NOTES:  You can use HOUR to select data for a reporting
period.


+-----------------+
| I N T E R V L S |   Number of Recording Intervals
+-----------------+
INTERVLS is the number of measurement intervals that occurred
over a specific measurement period (DURATION).  It is the
number of monitor logical records.

A data collector collects its measurements for a user-defined
period of time and then writes them to an output record.
Each such period of time is termed an "interval" and its
length, which extends from the record's starting timestamp
(STARTTS) to its ending timestamp (ENDTS), is termed its
duration.  The derivation is potentially unique to each of
the data collectors.

TYPE:  Common retained data element --  NUMERIC, measurement
unit is INTEGER.

SEE ALSO:  DURATION - Recording Interval Time
            STARTTS  - Start Time Stamp
            ENDTS    - End Time Stamp

ELEMENT DERIVATION:  For RMF, CICS, IMS, Network, TAP, and
MQSeries analyzer files in the DETAIL timespan, INTERVLS=1.

For the TLM component of the CA MICS Tape Analyzer Option,
INTERVLS is always one at the detail timespan.

For TSO/MON files in the DETAIL timespan, INTERVLS only
exists in the TSOTSO, TSOTSU, and TSOTSC files, which are all
derived from the TSO/MON Systems Record.

INTERVLS=1 for all records in the TSOTSO file in the DETAIL
timespan.

INTERVLS may be 1 or greater in the TSOTSU file.  CA MICS
combines user information from multiple TSO/MON Systems
Records when the user was completely inactive for one or more
TSO/MON recording intervals.  In this case, INTERVLS is equal
to the number of combined recording intervals.

Since the TSOTSC file is summarized to the hour in the DETAIL
timespan (necessary for data compression), INTERVLS is equal
to the number of recording intervals that occurred during
the hour (usually 12 at the recommended TSO/MON Systems
Record recording interval of 5 minutes).

For CA MICS Analyzer Option for VM/CMS files in the DETAIL
timespan:

In VMCVMU and all VMX interval recording files, INTERVLS is
computed by dividing DURATION of a record by the length of
recording interval.

In VMX event files and other VMC files in which INTERVLS
exists, INTERVLS is set to 1.

RANGE OF VALUE:  The value of this element ranges from a
minimum of 0 to a maximum that is limited only by practical
bounds.

SUMMARIZATION PROCESS:  INTERVLS contains the sum of all
occurrences of INTERVLS that were included in the level of
summarization.
```

| M O N T H | Month of Year

MONTH is the calendar month of the year (1 through 12).

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: DAY - Day of Month

ELEMENT DERIVATION: MONTH is derived from the standard CA MICS time-stamp field ENDTS. The value ranges from 1 through 12, representing January through December.

If the 13MONTHYEAR option is specified in MICS.PARMS(SITE), then the value of MONTH ranges from 1 to 13, representing the thirteen months in the organization's fiscal year. In this case, the fiscal year is divided into groups of 28 days, and MONTH is set according to which group of 28 days the date/time falls. For example, if the fiscal year was specified as starting on July 1, July 28 would be the last day of MONTH=01 and July 29 would be the first day of MONTH=02.

The MONTH variable in the CA MICS Accounting and Chargeback financial recap files, ACTRCP and ACT_RC, is set to the month value of the prior month. MONTH in these files represents the month for which the work was invoiced, not the month in which the work occurred.

RANGE OF VALUE: MONTH takes any value in the following ranges:

1 to 12; 1 to 13 with 13MONTHYEAR option.

SUMMARIZATION PROCESS: MONTH contains the last value of MONTH processed in the level of summarization.

At the higher levels of summarization in most CA MICS files, MONTH is not summarized. It is used to determine a unique key for records within a file based on the timespan, so other elements in the file can be summarized.

| O R G S Y S I D | Original System Identification

For z/OS systems, ORGSYSID is the SMF system identifier of the z/OS system on which data was collected.

The z/OS SMFWTM routine copies field SMCASID to offset x'0E' of every SMF record before it writes it to a SYS1.MANx data set. SMCASID is filled in during the IPL process from the SID= parameter in member SMFPRMxx of SYS1.PARMLIB.

For non-z/OS systems, ORGSYSID is not obtained directly from the input data. It is obtained from specifications in prefix.MICS.PARMS member cccOPS.

TYPE: Common retained data element - CHARACTER, LENGTH 4.

SEE ALSO: SYSID - System Identification

ELEMENT DERIVATION: When z/OS SMF is the source of the input records, ORGSYSID is extracted from the field at offset 10 in the standard SMF header area. In IBM documentation, this field has the name SMFxxSID.

Although the VCA component normally reads SMF records written by the VCC component, it will also read non-SMF SMF-look-a-like records written by VCC to a sequential data set.

For all AUM, DEX, VSE/PWR, and SRL files, ORGSYSID is not obtained directly from the data. It is obtained from specifications in prefix.MICS.PARMS member cccOPS.

For CICS, when reading data written by ASG-TMON for CICS, the ORGSYSID is equated to the field called xxSMFSID in the ASG-TMON for CICS documentation. The 'xx' stands for the record type: TA, TI, or TR.

For IMS, when reading data extracted from the BMC MAINVIEW for IMS Online product and the IMS Analyzer AutoDetect Facility is being used, the ORGSYSID value is extracted from the MAINVIEW transaction accounting records, specifically from the field containing the System Name ID (SYSNAME). Otherwise, the SYSID specified on the IMSOPS OPTION statement's OSYSID parameter is assigned to all the incoming log data.

NOTE: A more correct approach to populating the ORGSYSID value using the IMS Analyzer AutoDetect Facility would be to use the field that contains the SMF ID to populate the ORGSYSID data element. However, the BMC Mainview transaction accounting records do not always contain the SMF ID. If your System Name IDs (SYSNAME) differ from the SMF IDs defined for the z/OS LPARs, the IMS00408 error message displays in the MICS log with information about the SYSNAME and SMF ID values so that you can make appropriate adjustments to the OSYSID parameter of the IMSOPS OPTION statement.

Tandem MEASURE data contains a Tandem system identifier. The

CA MICS Tandem Option matches this Tandem system identifier to one of those specified in the TDMOPS parameter library member to identify the Tandem system and obtain the ORGSYSID.

For the TLM component of the CA MICS Tape Analyzer Option, ORGSYSID represents the tape library name (TAPELIB) used to identify unique tape library systems. ORGSYSID is derived from the TMCTLIB field of the TMCC record. TMCTLIB is the tape library name specified during TMCC execution. ORGSYSID is not used in the sort sequence of TLM files and is not retained outside the DETAIL timespan.

For the VTS component of the CA MICS Tape Analyzer Option, since no ORGSYSID is supplied by VTS data sources, ORGSYSID is not used. Data element SYSID is assigned the value SYS1 by default, unless overridden with a VTSSYSID statement in prefix.MICS.PARMS(VTSOPS).

VM measurement data does not contain the SMF System Identifier. The CA MICS Analyzer Option for VM/CMS matches the CPU model and serial number in the monitor data to an ORGSYSID definition in the VMCOPS parameter library member to identify a VM system and obtain the ORGSYSID.

For VSE/PWR files, this field is obtained from the VSESYSID based on the mapping defined in prefix.MICS.PARMS(PWROPS).

For the CA MICS Web Analyzer component processing Web Log Extractor (WLE) records, the ORGSYSID is supplied by configuration file server information parameters in prefix.MICS.PARMS(WLECONF).

For a non-MSI FDA component, ORGSYSID is populated by code in the USRFMT2 macro. The MSI FDA uses the ORGSYSID statement in its fdaGENIN member to identify the input field containing the ORGSYSID. The MSI FDA could also use the cccSYSID exit for specialized ORGSYSID assignment.

RANGE OF VALUE: 1 to 4 alphanumeric characters.

SUMMARIZATION PROCESS: ORGSYSID contains the value of the last ORGSYSID encountered.

USAGE NOTES: CA MICS provides a mapping of input SYSID (ORGSYSID) to logical SYSID (SYSID) in order to support the concept of "successor" systems.

The ORGSYSID is used to find information about the computer system that is identified by the ORGSYSID, information such as the CPU model and MIPS rating.

S T A R T T S | Start Timestamp

STARTTS has different definitions depending on the timespan and the data collector.

DETAIL Timespan:

- o Gives the time of initiation for an executable function, such as job or command.
- o Tells when a measurement interval started.

All Other Timespans:

- o Gives the minimum start timestamp encountered in the records used to compile the summarized record.
- o In any timespan but DETAIL, the combination of STARTTS and ENDTS in the record identifies the boundaries of the range of time over which the data in the record was summarized.

In the CA MICS Accounting and Chargeback Option ledger files, STARTTS is the lowest ENDTS value in the journal files from which the ledger file was summarized.

TYPE: Common, retained - Numeric, measurement unit is SAS datetime value.

SEE ALSO: DURATION - Recording Interval Time
ENDTS - End Timestamp

ELEMENT DERIVATION: The derivation of STARTTS depends on the file:

RMF Files: For files generated from the SMF type 70-79 RMF records, STARTTS is derived from fields SMFnnDAT (interval start date) and SMFnnIST (interval start time) of each RMF record (where nn is the record number).

For files generated from the SMF type 113 subtype 1 or 2 Hardware capacity, reporting, and statistics record generated by the IBM CPU Measurement Facility (CPU MF) Hardware Instrumentation Services (HIS).

For 113-2 records, STARTTS is determined in one of two ways:

- If the interval record is the first from a HIS collection run, STARTTS is set equal to field SMF113_2_CTS, adjusted to local time by adding the GMT Offset value.

- If the interval record is an intermediate or final from a HIS collection run, STARTTS is set equal to the ENDTS of the previous interval record.

For 113-1 records, STARTTS is set equal to field SMF113IntervalStart, adjusted to local time by adding the GMT Offset value.

- If the interval record is the first from a HIS collection run, STARTTS is set equal to field SMF113_2_CTS, adjusted to local time by adding the GMT Offset value.

- If the interval record is an intermediate or final from from a HIS collection run, STARTTS is set equal to the ENDTS of the previous interval record.

TSO Files: For the TSOTSO file, this value is calculated by subtracting the length of recording interval (TSMSINTV) from the record generation date (TSMSDTSP) and record generation time (TSMSTMSP) fields in the systems segment of the TSO/MON Systems Record.

For the TSOTSU and TSOTSC files, STARTTS is the same as either the TSOTSO file or the LOGON timestamp, TSMSLOGD (LOGON date) and TSMSLOGT (LOGON time), whichever is less.

For the TSOTSI file, STARTTS is taken from the Command Start Time (TSMPSTRT) and Command Start Date (TSMPSTRD) fields in the TSO/MON Command Record.

SMF Files: For the BATJOB file, STARTTS is taken from fields SMF30DTE and SMF30TME in the Common Address Space Work Record (type 30), subtype 5.

For the BATPGM, BAT_ST, BAT_TS, and BAT_TP files:

o For complete step records (PGMINTVL='SE24'), STARTTS is taken from fields SMF30STD (date initiator selected step) and SMF30SIT (time initiator selected step) of the SMF Common Address Space Work Record (type 30).

o For interval records, STARTTS represents the start time of the interval and is taken from one of the following, depending on the level of z/OS that produced the SMF record:

- The end time of the previous interval
- Fields SMF30IST and SMF30ISD
- Field SMF30ISS

For the BATSPL file, STARTTS is taken from fields SMF6WSD and SMF6WST (datetime output function began work on this file) of the SMF JES2 and JES3 Output Writer records (type 6).

ADM Files: For ADMEXC at the DAYS level, STARTTS is undefined. For ADMEXC at the MONTHS level, STARTTS is the lowest date and time associated with data summarized into the record.

CA IDMS Files: For the IDMS System Incident File (IDMSIN), STARTTS is the date and time when the CA IDMS transaction began processing. STARTTS contains the value of the Start Time (PMHSTIME) field from the header section of the CA IDMS Performance Monitor record.

For all other CA IDMS files, STARTTS is adjusted based on a user-defined interval value. This value is coded in the INTERVAL statement of prefix.MICS.PARMS(IDMOPS) and determines the cut-off point for each timespan.

CICS Files: For the CICS Incident File (CICCIN), STARTTS is set to the value of the ENDTS for the incident. The STARTTS value for the Application Unit Activity File (CICCAU) is set from and identical to the STARTTS value for the User Activity File (CICCSU). Construction of the timestamp itself varies according to the monitor used to collect CICS data:

CMF: When type 110 Performance records are input, the STARTTS is read from Field 5 (START) that carries either the time at which the user task was attached or the time at which data recording was most recently reset in support of the MCT user event monitoring point DELIVER option or of the monitoring options MNCONV, MNSYNC, or FREQUENCY. After this value is read from Field 5, it is adjusted by the GMT offset that is input from the header section of the same type 110 record. The resulting value populates the STARTTS in the CICS User Activity (CICCSU) file and in the CICS Application Unit Activity (CICCAU) file.

When type 110 Statistics records are input, the STARTTS is calculated by the formula:

ENDTS - DURATION

The STARTTS values in this case populate the CICS System Activity (CICCSY) file.

ASG-TMON for CICS TS: To populate the CICS User

Activity (CICCSU) file, the STARTTS is read from the main segment of the type TA record. To populate the CICS System Activity (CICCSY) file, the STARTTS is read from the main segment of the type TI record.

DB2 Files: For the DB2DSU, DB2DSP, DB2DDU, and DB2DDP files, STARTTS is set equal to ENDTS minus the elapsed time, DSUELPTM. For the DB2DSD, DB2DSY, and DB2DDY files, STARTTS is set equal to the ENDTS of the previous record.

IMS Files: For the IMS Incident File (IMSINC), STARTTS is equal to and set from the ENDTS of the incident. The STARTTS values for the Application Unit Activity File (IMSIAU) and User Application Count File (IMSIAC) are set from and identical to the STARTTS value for the User Activity File (IMSISU). For the System Activity File (IMSISY), STARTTS = ENDTS - DURATION. Construction of STARTTS for transaction-oriented files varies according to the method used to accumulate IMS data:

With IMF: STARTTS is constructed from the time and date of processing region occupancy presented by Control/IMS in its detail transaction record.

Without IMF: STARTTS is constructed from the time and date fields in the log type '31'X record for the Get Unique of the transaction's first message segment from the IMS message queue. The field in the '31'X record that contains date and time is QLGUDTME.

MQS Files: For all files except MQAQAA, MQATAA, and MQBMFA, the STARTTS for an interval is the ENDTS of the previous interval.

For the MQAQAA and MQATAA files, STARTTS is derived by subtracting Task Execution Time (TAAEXCTM) from End Timestamp (ENDTS): STARTTS = ENDTS - TAAEXCTM;

Note that TAAEXCTM = TAAIETS-TAAISTS where TAAIETS is SMF field WTASINTE and TAAISTS is SMF field WTASINTS in IBM DSECT CSQDWTAS.

For the MQBMFA file, STARTTS is set to the interval start date and time values found in the Message flow accounting record fields IMFLSTDT and IMFLSTTM.

For Web Analyzer:

STARTTS is set to the value input from field SM1209CM of the SMF type 120 record subtype 9 at offset 4 (x'04') of the z/OS Request Information section.

Network Analyzer Option Files:

NPM Network data collection: STARTTS is the end timestamp of the measurement interval minus the actual observation time of the interval: STARTTS = ENDTS - NPMTMAOT.

NPM Session data collection: STARTTS is the Date and Time of Beginning Interval, VPVDT, from the Session Statistics Detail Record.

NetView/NLDM or CA NetMaster NTS data (by CA MICS File):

- o NVSRTM and SNTNSV files - STARTTS is taken from field LRTMCOLB of the SMF type 39 record.
- o NVSNSA file - STARTTS is taken from field LACCBEGT of the Accounting and Availability data section of SMF type 39 record.
- o NVSNSC and NVSRTE files - There is only one observation in CA MICS for each unique SNA session in the DETAIL timespan of these files, and the STARTTS for that observation is taken from the lowest timestamp found in all SMF type 39 records processed for that session.

NetView/NPDA or CA NetMaster NEWS Data: STARTTS is set equal to ENDTS.

IBM Sterling Connect:Direct data collection: STARTTS is the timestamp for when the COPY function was started. It is a combination of fields CTSDATE and CTSTIME from the SMF 133 (user defined) record, subtype CT.

VSE/Power Files: The element is derived from the ACDATE and ACSTRT fields of the account record used to create the observation. However, if the value of ACSTRT is greater than ACSTOP, the date used is one day earlier than ACDATE.

When processing the Execution Account Record for VSE/Power, the element is derived by subtracting the execution time (EXCJDUR) from ENDTS.

VM/CMS Files: In the DETAIL timespan, the derivation of STARTTS varies on files:

For VMCCMU, VMCDAC, and VMCVRA, STARTTS contains the difference between ENDTS and DURATION. In VMCNTU file, STARTTS contains the value input from record type 07 column 45-52.

For VMXVXU, VMXVXT, VMX_VU, and VMX_VT, STARTTS of a user record is set to the value of ENDTS of the previous record processed for the user. For the first user record after a user logged on, STARTTS is set to either the logon timestamp, if available, or the difference between ENDTS and interval record length.

For other files, STARTTS contains the value of ENDTS of the previous record processed under the same unique key for records within a file.

TLM Component TMC Information Area files: For the TMCfff files, STARTTS is the start of the collection interval. Generally, this is the last time TMCC was executed for a given tape library. STARTTS is derived from the TMCLDATE/TMCLTIME fields of the TMCC record.

VTS Component IVT Information Area files: For the IVTfff files, created from SMF type 94 records for IBM VTS systems, STARTTS is established as the exact hour value of the hour that preceded the hour when the SMF type 94 record was written to the SMF buffer.

For example, if an SMF type 94 record was written at 13:12:48.02, the STARTTS for all IVTfff files generated from that record would be 12:00:00.00.

VTS Component SVT Information Area files: For the SVTSSP, SVTCIP, and SVTRPR files, STARTTS is created from Sun StorageTek User SMF records. It is calculated by subtracting 15 minutes from the ENDTS for the interval records, subtypes 10, 11, and 20.

For all other SVT files except SVTSVS, STARTTS is set to the ENDTS value read from the StorageTek event records.

For the SVTSVS file, STARTTS is set to the start of each 15 minute boundary 00:00.00, 15:00.00, 30:00.00, 45:00.00, in the DETAIL timespan.

RANGE OF VALUE: Any valid SAS datetime value. The value of this data element is not affected by your use of the 13MONTHYEAR option in prefix.MICS.PARMS(SITE).

SUMMARIZATION PROCESS: The minimum STARTTS encountered in summarizing a number of records to a higher level of granularity.

| S Y S I D | SYSTEM IDENTIFICATION

SYSID is the translated system identifier that was derived from ORGSYSID by the CA MICS SYSID translation process. It is defined in the CPLXSID or SYSID member of a CA MICS parameter library.

TYPE: Common, retained - Character, length 4.

SEE ALSO: ORGSYSID - Original System Identifier

ELEMENT DERIVATION: Every record coming from most CA MICS data collection sources contains an SMF system identifier in the appropriate field of the record. CA MICS calls this field the ORGSYSID. For collection sources that do not provide such an identifier, CA MICS provides a method to determine the ORGSYSID. This ORGSYSID is compared to a user-defined SYSID table (PARMS members CPLXSID and/or SYSID) to determine if the ORGSYSID should be used as the SYSID or translated to a reporting SYSID.

The prefix.MICS.PARMS(SYSID) member is used to build SYSID table entries for a specific unit. The sharedprefix.MICS.PARMS(CPLXSID) member is used to build SYSID table entries for every unit in a complex.

RANGE OF VALUE: SYSID can contain any combination of 1 to 4 alphabetic and/or numeric characters.

SUMMARIZATION PROCESS: SYSID is not summarized. It is used to determine a unique key for records within a file based on the timespan, so other elements in the file can be summarized.

USAGE NOTES: Most CA MICS files have SYSID as their highest level or primary key. This is what allows CA MICS to store information from multiple systems (SYSIDs) in the same database.

The ADMIHL file uses ORGSYSID as its highest level or primary key.

The CA MICS Network Analyzer Option uses SYSID as its second highest level key. The highest level key in these files is NETWRKID, Network Identifier. NETWRKID enables these files to support data from multiple networks (for example, networks supported by SNA Network Interconnect, SNI).

For the VTS component of the CA MICS Tape Analyzer Option, in all IBM Virtual Tape Subsystem Information Area (IVTfff) files, SYSID defaults to SYS1, or may be set to a value

specified in prefix.MICS.PARMS(VTSOPS) using the VTSSYSID statement.

For the TLM component of the CA MICS Tape Analyzer Option, SYSID is not used in the sort sequence of TLM files, and is not retained outside the DETAIL timespan.

There is no general way of meaningfully combining or even comparing data gathered on different computer systems, so be careful when putting together data from two SYSIDs.

W E E K	WEEK OF YEAR

WEEK is the calendar week of the year (an integer from 1 to 54). To CA MICS, a week extends from Sunday through Saturday (unless changed via the WEEKSTART parameter mentioned below). The first and last weeks of the year may have less than seven days.

WEEK=54 can only occur in a leap year, and then only when January 1 falls on a Saturday (and therefore is a week all by itself).

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: DAY - Day of Month
MONTH - Month of Year

ELEMENT DERIVATION: WEEK is calculated from the date portion of the variable ENDTS. WEEK=01 begins on January 1 and extends through the following Saturday. WEEK=02 begins on Sunday, and so forth.

If you specified the 13MONTHYEAR option of MICS.PARMS(SITE), then WEEK=01 begins on the first day of the fiscal year specified and extends up to the end of the following Saturday. If you have modified the definition of a standard week via the WEEKSTART parameter in MICS.PARMS(SITE), then your WEEK=01 will begin on the first day of your fiscal year and extend to the end of the day before that specified at the operand of the WEEKSTART statement.

RANGE OF VALUE: WEEK takes any numeric value in the range 1 through 54.

SUMMARIZATION PROCESS: WEEK is set to the value of the last WEEK encountered in the summarization period.

For files that exist in the WEEKS timespan, WEEK is not summarized. It is used to determine a unique key for records in the WEEKS timespan so other elements in the file can be summarized.

Y E A R and Y E A R 4	Year of Century

The YEAR (the encoded year of the century) and YEAR4 (the four digit year) data element descriptions are presented in the following order:

* YEAR -- Year of Century
* YEAR4 -- Year

Y E A R	Year of Century

YEAR is the "cyy" encoded year of the century where yy is the last two-digits of year and c, the century indicator, is 0 for years before 2000 and 1 beginning in year 2000. For example, the year 1998 is represented as 98 while the year 2001 is represented as 101.

You can also think of YEAR as the 4-digit year minus 1900.

YEAR is carried in encoded form to reduce database space while preserving proper chronological sort sequence (e.g., data for year 2000 sorts after data from 1999). Since YEAR appears in every observation of every file, a savings of several megabytes is possible. The maximum value of the encoded YEAR element is 255, which supports years through 2155 (1900 + 255).

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: MONTH - Month of Year
WEEK - Week of Year

ELEMENT DERIVATION: Normally, CA MICS calculates YEAR by subtracting 1900 from the year part of ENDTS.

If 13MONTHYEAR YES was specified in prefix.MICS.PARMS(SITE), YEAR is set to the year in which the fiscal year defined by the time stamp begins minus 1900. In other words, if your site specifies that your fiscal year runs from July 1 of one year to June 30 of the following year, then YEAR would be set to 201 for any date/time from 00:00:00 on July 1, 2001, to 23:59:59 on June 30, 2002.

The YEAR variable in the CA MICS Accounting and Chargeback

financial recap file is set to the year value of the prior period. In all of the financial recap files: ACTRCP, ACT_RC, ACT_YR, ACT_YT, ACT_RR, AND ACT_RT, the variable YEAR represents the year for which the work was invoiced.

RANGE OF VALUE: YEAR takes any numeric value in the range 0 to 255.

SUMMARIZATION PROCESS: YEAR contains the value of the last YEAR encountered during the summarization process.

At the higher levels of summarization in most CA MICS files, this is a summarization control variable. During the summarization process a new record is created for each unique record key combination, and this field is part of the key of its record.

USAGE NOTES: The default output format for YEAR, YR., prints the last 2-digits of year to retain the "year of century" presentation. For example, for data from year 2001, YEAR is printed as 01 rather than the actual value of 101.

For some applications, such as in program logic, it may be helpful to work with the full 4-digit year. The %YEAR4 macro is available for this purpose. For example, the following SAS instructions both select data for the years 1998 through 2001.

```
IF 98 LE YEAR LE 101;

IF 1998 LE %YEAR4 LE 2001;
```

Y E A R 4	Year

YEAR4 is the four digit year.

TYPE: Common, retained - Numeric, measurement unit is integer.

SEE ALSO: MONTH - Month of Year
WEEK - Week of Year
YEAR - Year of Century

ELEMENT DERIVATION: CA MICS calculates YEAR4 as the value of the element YEAR plus 1900.

RANGE OF VALUE: YEAR4 takes any numeric value in the range 0 to 2155.

SUMMARIZATION PROCESS: This value is not summarized.

USAGE NOTES: The YEAR4 data element is shipped turned off. Furthermore, it is not defined as a normal CA MICS data element in sharedprefix.MICS.GENLIB. As a result, you cannot use the standard method of file tailoring to turn YEAR4 on or off. You must activate or deactivate this data element via a parameter in sharedprefix.MICS.PARMS(CPLXDEF). Since CPLXDEF affects the operation of the entire CA MICS complex, the data element will either reside in every CA MICS file or none at all.

Refer to section 2.3.1.8 of the "PIOM" for information on activating/deactivating the YEAR4 data element.

Z O N E	Time Zone

ZONE is the time zone during which the measurements in this record were taken.

ZONE enables you to categorize activity at a level of granularity that is greater than HOUR but less than DAY. (This is in the DETAIL timespan; ZONE can contain much more time than a day in the higher timespans.)

Each installation defines ZONEs in the manner that is most meaningful to it, but most installations include the notion of "prime" vs. "non-prime" time in the ZONE definitions.

The definition and use of ZONEs is discussed in Section 2.3.2.3 of the CA MICS Planning, Installation, Operation, and Maintenance (PIOM) Guide.

TYPE: Common, retained - Character, length 1.

SEE ALSO: HOUR - Hour of Day

ELEMENT DERIVATION: CA MICS derives ZONE by comparing the date and time (ENDTS) of the measurement source record with your site's HOLIDAY specifications. If there is a match, the "holiday" ZONE is assigned. If there is not, CA MICS compares the DAY and HOUR with the your ZONE specifications in member ZONE of prefix.MICS.PARMS. A match will always be found because MICSGEN requires that all 168 hours within a week be assigned to a ZONE.

RANGE OF VALUE: ZONE contains any one of the following characters:

1, 2, 3, 4, 5, 6, 7, 8, or 9.

SUMMARIZATION PROCESS: ZONE is set to the value of the last ZONE encountered in the summarization process.

At the higher levels of summarization in most CA MICS files, ZONE is not summarized. It is used to determine a unique key for records within a file based on the timespan, so other elements in the file can be summarized.

USAGE NOTES: ZONE is not affected by the 13MONTHYEAR option in prefix.MICS.PARMS(SITE).

B.2 - The UNIX Global Performance (AUMGPR) File

The UNIX Global Performance file contains global system measurements useful for capacity, system load, and performance analyses. It includes information about CPU use, disk I/Os, paging and swapping, and wait queues, and is derived from MeasureWare Global metrics.

The metrics supported vary by operating system.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity
DETAIL	SYSID AUMOPSYS YEAR MONTH DAY HOUR ENDTS
DAYS	SYSID AUMOPSYS YEAR MONTH DAY HOUR
WEEKS	SYSID AUMOPSYS YEAR WEEK ZONE HOUR
MONTHS	SYSID AUMOPSYS YEAR MONTH ZONE
YEARS	SYSID AUMOPSYS YEAR ZONE
TABLES	N/A

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
	Sequence/Summary Data Elements	
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	SYSID	- System Identifier
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
	Common Data Elements	
XDWMY.E	AUMCPMDL	- System Model Identification
XDWMY.E	AUMSAMPL	- Number of Data Samples
XDWMY.E	AUMSYSID	- Name of Host System
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
	Retained Data Elements	
XDWMY.E	GPRACCPU	- CPUs Online per System
XDWMY.E	GPRSUPTM	- Time of Last System Reboot
	Accumulated Data Elements	
XDWMY.E	GPRACTVM	- Active Virtual Memory Size
XDWMY.E	GPRAPRCT	- Sum of Avg Processes Active
XDWMY.E	GPRBIOCT	- Processes Blocked on Local Disks
XDWMY.E	GPRCACHE	- Size of Buffer Cache Configured
XDWMY.E	GPRCHUSE	- The Size of Currently Used Buffers
XDWMY.E	GPRCPQTM	- CPU Queue Wait Time
XDWMY.E	GPRCPRCT	- Total Processes Completed
XDWMY.E	GPRCPUCP	- Total Available CPU Time
XDWMY.E	GPRCPUTM	- Total CPU Busy
XDWMY.E	GPRCSWTM	- CPU Busy for Context Switching
XDWMY.E	GPRDBKCT	- Block I/Os
XDWMY.E	GPRDBRCT	- Block Reads
XDWMY.E	GPRDBWCT	- Block Writes
XDWMY.E	GPRDFRCT	- Disk Reads by File System
XDWMY.E	GPRDFWCT	- Disk Writes by File System
XDWMY.E	GPRDKQTM	- Disk I/O Queue Wait Time
XDWMY.E	GPRDMMCT	- Virtual Memory Management I/Os
XDWMY.E	GPRDMRCT	- Disk Reads for Virt Memory Management
XDWMY.E	GPRDMWCT	- Disk Writes for Virt Memory Management
XDWMY.E	GPRDRRCT	- Raw Reads
XDWMY.E	GPRDRWCT	- Raw Writes
XDWMY.E	GPRDSKIO	- Physical Disk I/Os
XDWMY.E	GPRDSRCT	- Disk Reads for File System Management
XDWMY.E	GPRDSWCT	- Disk Writes for File System Management
XDWMY.E	GPRFSDCT	- File System I/Os
XDWMY.E	GPRFSMCT	- File System Management I/Os
XDWMY.E	GPRINODU	- Inode Entries in Use
XDWMY.E	GPRINTTM	- CPU Busy for Interrupt Handling
XDWMY.E	GPRIPQTM	- IPC Queue Wait Time
XDWMY.E	GPRLCOLS	- LAN Collisions
XDWMY.E	GPRLDKCT	- Logical Disk I/Os
XDWMY.E	GPRLERRS	- LAN Packet Errors
XDWMY.E	GPRLOSTB	- Trace Buffers Lost
XDWMY.E	GPRLPCKI	- Successful LAN Packets Received
XDWMY.E	GPRLPCKO	- Successful LAN Packets Sent

```
XDWMY.E  GPRLRDCT - Logical Disk Reads
XDWMY.E  GPRLRDKB - KBytes Logically Read From Disk
XDWMY.E  GPRLUSED - Number of File Locks in Use
XDWMY.E  GPRLWTCT - Logical Disk Writes
XDWMY.E  GPRLWTKB - KBytes Logically Written To Disk
XDWMY.E  GPRMEMSW - Amount of Swap Space Configured
XDWMY.E  GPRMEMSZ - Amount of Main Memory Configured
XDWMY.E  GPRMEMUS - Amount of Main Memory Available for User
XDWMY.E  GPRMEMUT - Main Memory Used
XDWMY.E  GPRMMEMS - Main Memory Used by System
XDWMY.E  GPRMMEMU - Main Memory Used by Users
XDWMY.E  GPRMMQTM - Memory Queue Wait Time
XDWMY.E  GPRMQUSE - Message Queues in Use
XDWMY.E  GPRMSGQU - Message Queues Available
XDWMY.E  GPRNFSCT - Network File System Requests
XDWMY.E  GPRNICTM - CPU Busy at Nice Priorities
XDWMY.E  GPRNORTM - CPU Busy at Normal Priorities
XDWMY.E  GPRNPROC - Value of NPROC Parameter
XDWMY.E  GPRNTQTM - Network Queue Wait Time
XDWMY.E  GPROFLOW - Unmeasurable New Transactions
XDWMY.E  GPROUTQU - Outbound Queue Lengths
XDWMY.E  GPRPAUTM - CPU Idle Due to I/O Wait
XDWMY.E  GPRPFLCT - Page Faults
XDWMY.E  GPRPGOCT - Number of Page Outs
XDWMY.E  GPRPKBCT - KBytes Physically Transferred
XDWMY.E  GPRPRDCT - Physical Disk Reads
XDWMY.E  GPRPRITM - Time Blocked on Priority
XDWMY.E  GPRPROCT - Sum of Avg Processes Alive
XDWMY.E  GPRPTBLU - Proc Table Entries Used
XDWMY.E  GPRPWTCT - Physical Disk Writes
XDWMY.E  GPRRAWCT - Raw I/Os
XDWMY.E  GPRRDHIT - Buffered Read Hits
XDWMY.E  GPRRTMTM - CPU Busy at Real Time Priorities
XDWMY.E  GPRRUNTM - Total Run Time of Completing Processes
XDWMY.E  GPRSCALL - Total System Calls
XDWMY.E  GPRSEMID - Semaphore Identifiers Available
XDWMY.E  GPRSEMST - Semaphore Identifiers in Use
XDWMY.E  GPRSHMSG - Table Memory Segments Available
XDWMY.E  GPRSLQTM - Sleep Queue Wait Time
XDWMY.E  GPRSMEMS - Shared Memory Segments in Use
XDWMY.E  GPRSPRCT - Total Processes Started
XDWMY.E  GPRSSCTM - CPU Busy for System Calls
XDWMY.E  GPRSUSCT - Procs Suspended During Interval
XDWMY.E  GPRSWACT - Processes Waiting to be Swapped in
XDWMY.E  GPRSWAPU - Swap Space Used
XDWMY.E  GPRSWOCT - Number of Swap Outs/Deactivations
XDWMY.E  GPRSWPCT - Process Swaps, Activation/Deactivation
XDWMY.E  GPRSYSTM - CPU Busy in System Mode
XDWMY.E  GPRTBLEN - File Table Entries Used
XDWMY.E  GPRTBLUK - Entries Used by Kernel
XDWMY.E  GPRUSRCT - Sum of Avg Login Users
XDWMY.E  GPRUSRTM - CPU Busy in User Mode
```

Maximum Data Elements

```
XDWMY.E  GPRMPBDU - Percent Utilization of Busiest Disk
XDWMY.E  GPRMPDSU - Peak Disk Space Utilization
XDWMY.E  GPRMXCPU - Max Number of Processors Configured
XDWMY.E  GPRMXDSK - Max Number of Disk Drives Configured
XDWMY.E  GPRMXDTM - Time Busiest Disk Was Busy
XDWMY.E  GPRMXLAN - Max Number of LAN Interfaces Configured
```

Derived Data Elements

```
XDWMY.E  GPRAVAPR - Avg Number of Processes Active
XDWMY.E  GPRAVBIQ - Avg Processes Blocked on Local Disks
XDWMY.E  GPRAVCPQ - Avg Processes in Run Queue
XDWMY.E  GPRAVDKQ - Avg Processes Waiting for Disk I/Os
XDWMY.E  GPRAVIPQ - Avg Processes Waiting for IPC Resources
XDWMY.E  GPRAVLCL - Rate of LAN Collisions per Minute
XDWMY.E  GPRAVLER - Rate of LAN Errors per Minute
XDWMY.E  GPRAVMMQ - Avg Processes Waiting for Memory
XDWMY.E  GPRAVNTQ - Avg Processes Waiting for Network
XDWMY.E  GPRAVPRO - Avg Number of Processes Alive
XDWMY.E  GPRAVPRQ - Avg Processes Blocked on Priority
XDWMY.E  GPRAVRTM - Avg Run Time of Completing Processes
XDWMY.E  GPRAVSLQ - Avg Processes in Sleep State
XDWMY.E  GPRAVSUS - Avg Number of Processes Suspended
XDWMY.E  GPRAVSWA - Avg Number of Processes Waiting Swap-In
XDWMY.E  GPRAVUSR - Avg Number of Login Users
XDWMY.E  GPRDKLRT - Logical Disk I/O Rate
XDWMY.E  GPRDMMRT - Virtual Memory Management I/O Rate
XDWMY.E  GPRDMRRT - Virt Memory Management Disk Read Rate
XDWMY.E  GPRDMWRT - Virt Memory Management Disk Write Rate
XDWMY.E  GPRDSKRT - Physical Disk I/O Rate
XDWMY.E  GPRNFSRT - Rate of Network File System Requests
XDWMY.E  GPRPAGRT - Page Fault Rate
XDWMY.E  GPRPCAVM - Percent of Virtual Memory in Active Use
XDWMY.E  GPRPCCPU - Percent Total CPU Busy
XDWMY.E  GPRPCCSW - Percent CPU Busy for Context Switching
XDWMY.E  GPRPCHIT - Cache Hit Ratio
XDWMY.E  GPRPCINT - Percent CPU Busy for Interrupt Handling
XDWMY.E  GPRPCMEM - Percent of Main Memory Used
XDWMY.E  GPRPCMMS - Percent of Main Memory Used by System
XDWMY.E  GPRPCMMU - Percent of Main Memory Used by Users
XDWMY.E  GPRPCNIC - Percent CPU Busy at Nice Priorities
XDWMY.E  GPRPCNOR - Percent CPU Busy at Normal Priorities
XDWMY.E  GPRPCRTM - Percent CPU Busy at Real Time Priorities
XDWMY.E  GPRPCSSC - Percent CPU Busy for System Calls
XDWMY.E  GPRPCSWS - Percent of Swap Space Used
XDWMY.E  GPRPCSYS - Percent CPU Busy in System Mode
XDWMY.E  GPRPCUSR - Percent CPU Busy in User Mode
XDWMY.E  GPRSCLRT - System Procedure Call Rate
XDWMY.E  GPRSWORT - Swap-Out Rate
```

XDWMY.E GPRSWPRT - Swapping Rate

| G P R A C C P U | CPUs Online per System

GPRACCPU contains the number of CPUs online for the system represented by the System Identification (SYSID) data element.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPRACCPU is set to the value of the MeasureWare Global metric GBL_ACTIVE_CPU.

RANGE OF VALUE: GPRACCPU ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRACCPU contains the last value of GPRACCPU processed in the level of summarization.

USAGE NOTES: GPRACCPU contains the actual number of active processors for the system at the end of each global interval (usually 5 minutes). This provides up-to-date CPU configuration data for customers using vPAR dynamic processor reconfiguration.

| G P R A C T V M | Active Virtual Memory Size

GPRACTVM contains the number of kilobytes of active virtual memory space.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRPCAVM - Percent of Virtual Memory in Active Use
GPRMEMSW - Amount of Swap Space Configured

ELEMENT DERIVATION: At the DETAIL time-span, GPRACTVM is computed using MeasureWare Agent metrics in the following calculation:

```
GBL_MEM_ACTIVE_VIRT_UTIL*GBL_SWAP_SPACE_AVAIL_KB/100
```

where GBL_MEM_ACTIVE_VIRT_UTIL is from Global data and GBL_SWAP_SPACE_AVAIL_KB is from Configuration data.

RANGE OF VALUE: GPRACTVM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRACTVM contains the sum of all values of GPRACTVM processed in the level of summarization.

USAGE NOTES: None

| G P R A P R C T | Sum of Avg Processes Active

GPRAPRCT is the accumulation of the average number of active processes during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: AUMSAMPL - Number of Data Samples
GPRAVAPR - Avg Number of Processes Active

ELEMENT DERIVATION: At the DETAIL time-span, GPRAPRCT is computed using MeasureWare Global metrics in the following calculation:

```
GBL_ACTIVE_PROC*GBL_PROC_SAMPLE
```

RANGE OF VALUE: GPRAPRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRAPRCT contains the sum of all values of GPRAPRCT processed in the level of summarization.

USAGE NOTES: None

| G P R A V A P R | Avg Number of Processes Active

GPRAVAPR contains the average number of processes active (consuming CPU time) during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: AUMSAMPL - Number of Data Samples
GPRAPRCT - Sum of Avg Processes Active

ELEMENT DERIVATION: GPRAVAPR is computed using the following algorithm:

```
GPRAVAPR = AUMSAMPL;
IF AUMSAMPL
```

```
    THEN GPRAVAPR = GPRAPRCT/AUMSAMPL;
```

RANGE OF VALUE: GPRAVAPR takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRAVAPR contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: Processes that are alive for a portion of the interval and that consume CPU during that time are counted according to the percent of the interval that they are alive. For example, if a process ran for 1 minute during a 5 minute interval and consumed CPU during its life, it will be counted as one-fifth of an active process.

GPRAVAPR is a good overall indicator of the system's workload. A large number of active processes could indicate a CPU bottleneck. To determine if the CPU is a bottleneck, compare GPRAVAPR with GPRCPUTM (Total CPU Busy) and GPRAVCPQ (Avg Processes in Run Queue). If GPRCPUTM is near 100 percent and GPRAVCPQ is greater than one, there is a CPU bottleneck.

```
+-----------------+
| G P R A V B I Q |  Avg Processes Blocked on Local Disks
+-----------------+
```

GPRAVBIQ contains the average number of processes blocked on local disk resources such as IO or paging during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

```
SEE ALSO:  AUMSAMPL - Number of Data Samples
           GPRBIOCT - Processes Blocked on Local Disks
           GPRMPBDU - Percent Utilization of Busiest Disk
```

ELEMENT DERIVATION: GPRAVBIQ is computed using the following algorithm:

```
  GPRAVBIQ = AUMSAMPL;
  IF AUMSAMPL
    THEN GPRAVBIQ = GPRAPRCT/AUMSAMPL;
```

RANGE OF VALUE: GPRAVBIQ takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRAVBIQ contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: GPRAVBIQ is an indicator of disk contention among active processes. Normally it should be a very small number. To determine if the disk is a bottleneck, compare GPRMPBDU (Percent Utilization of Busiest Disk) with GPRAVBIQ (Avg Processes Blocked on Local Disks). If GPRMPBDU is near 100 percent and GPRAVBIQ is greater than one, a disk bottleneck is probable

```
+-----------------+
| G P R A V C P Q |  Avg Processes in Run Queue
+-----------------+
```

GPRAVCPQ contains the average number of processes in the run queue during the measurement interval. GPRAVCPQ measures the time as an average of total execution time for all processes. Processes marked "runable" include a process using the CPU, a process waiting for the CPU, and a process paused on a "short disk wait" at the time of the sample.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

```
SEE ALSO:  DURATION - Recording Interval Time
           GPRAVPRO - Avg Number of Processes Alive
           GPRCPQTM - CPU Queue Wait Time
```

ELEMENT DERIVATION: GPRAVCPQ is computed using the following algorithm:

```
  GPRAVCPQ = DURATION;
  IF DURATION
    THEN GPRAVCPQ = GPRCPQTM/DURATION;
```

RANGE OF VALUE: GPRAVCPQ takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRAVCPQ contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: It indicates a CPU bottleneck if GPRPCCPU (Percent Total CPU Busy) is 100 percent and GPRAVCPQ is greater than one.

The Analyzer's Daily Global Bottlenecks report helps you identify potential CPU, memory, and disk bottlenecks.

```
+-----------------+
| G P R A V D K Q |  Avg Processes Waiting for Disk I/Os
+-----------------+
GPRAVDKQ contains the average number of processes waiting for
disk and other non-terminal I/O during the measurement
interval.  This element measures the time processes were
blocked on disk, inode, CDFS, or general I/O waits as an
average of total execution time for all processes.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           GPRAVPRO - Avg Number of Processes Alive
           GPRDKQTM - Disk I/O Queue Wait Time

ELEMENT DERIVATION: GPRAVDKQ is computed using the following
algorithm:

  GPRAVDKQ = DURATION;
  IF DURATION
    THEN GPRAVDKQ = GPRDKQTM/DURATION;

RANGE OF VALUE:  GPRAVDKQ takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVDKQ contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  GPRAVDKQ is an indicator of disk contention
among active processes.


+-----------------+
| G P R A V I P Q |  Avg Processes Waiting for IPC Resources
+-----------------+
GPRAVIPQ contains the average number of processes waiting for
InterProcess Communications during the measurement interval.
The element measures processes that are blocked waiting for
completion of semaphore, message queue, pipe, socket, or
other IPC operations as an average of the total execution
time for all processes.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRIPQTM - IPC Queue Wait Time
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  GPRAVIPQ is computed using the following
algorithm:

  GPRAVIPQ = DURATION;
  IF DURATION
    THEN GPRAVIPQ = GPRIPQTM/DURATION;

RANGE OF VALUE:  GPRAVIPQ takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVIPQ contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  A large fluctuation in GPRAVIPQ from the normal
can indicate resource contention among processes using IPC.


+-----------------+
| G P R A V L C L |  Rate of LAN Collisions per Minute
+-----------------+
GPRAVLCL contains the number of LAN collisions detected, per
minute, during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           GPRLCOLS - LAN Collisions

ELEMENT DERIVATION:  GPRAVLCL is computed using the following
algorithm:

  GPRAVLCL = DURATION;
  IF DURATION
    THEN GPRAVLCL = 60*GPRLCOLS/DURATION;

RANGE OF VALUE:  GPRAVLCL takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVLCL contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  Collisions occur on any busy LAN, but abnormal
collision rates could indicate that too many packets are
being generated due to a hardware or software problem.
```

```
+-----------------+
| G P R A V L E R |  Rate of LAN Errors per Minute
+-----------------+
GPRAVLER contains the average number of LAN errors that
occurred per minute during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           GPRLERRS - LAN Packet Errors

ELEMENT DERIVATION: GPRAVLER is computed using the following
algorithm:

  GPRAVLER = DURATION;
  IF DURATION
    THEN GPRAVLER = 60*GPRLERRS/DURATION;

RANGE OF VALUE:  GPRAVLER takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVLER contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  A large error rate can indicate a hardware or
software problem.


+-----------------+
| G P R A V M M Q |  Avg Processes Waiting for Memory
+-----------------+
GPRAVMMQ contains the average number of processes that are
waiting for virtual memory management tasks during the
measurement interval.  This element measures the time
processes are blocked waiting for memory buffers or virtual
memory I/O (VM, cache, and MBUF) as an average of total
execution time for all processes.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRMMQTM - Memory Queue Wait Time
           DURATION - Recording Interval Time

ELEMENT DERIVATION: GPRAVMMQ is computed using the following
algorithm:

  GPRAVMMQ = DURATION;
  IF DURATION
    THEN GPRAVMMQ = GPRMMQTM/DURATION;

RANGE OF VALUE:  GPRAVMMQ takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVMMQ contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  GPRAVMMQ is an indicator of cache contention or
a memory bottleneck.


+-----------------+
| G P R A V N T Q |  Avg Processes Waiting for Network
+-----------------+
GPRAVNTQ contains the average number of processes waiting for
network resources such as LAN, NFS (Network File System), DUX
(diskless), RFA (Remote File Access) and RPC (Remote
Procedure Call) to become available during the measurement
interval.  This element measures the wait time as an average
of total execution time for all processes.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRNTQTM - Network Queue Wait Time
           DURATION - Recording Interval Time

ELEMENT DERIVATION: GPRAVNTQ is computed using the following
algorithm:

  GPRAVNTQ = DURATION;
  IF DURATION
    THEN GPRAVNTQ = GPRNTQTM/DURATION;

RANGE OF VALUE:  GPRAVNTQ takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVNTQ contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  GPRAVNTQ is an indicator of network bottleneck.
```

```
+-----------------+
| G P R A V P R O |  Avg Number of Processes Alive
+-----------------+
GPRAVPRO contains the average number of processes alive
during the measurement interval.  A process that was alive
for half of the measurement interval counts as half a
process.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  AUMSAMPL - Number of Data Samples
           GPRPROCT - Sum of Avg Processes Alive

ELEMENT DERIVATION: GPRAVPRO is computed using the following
algorithm:

  GPRAVPRO = AUMSAMPL;
  IF AUMSAMPL
    THEN GPRAVPRO = GPRPROCT/AUMSAMPL;

RANGE OF VALUE:  GPRAVPRO takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVPRO contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  None



+-----------------+
| G P R A V P R Q |  Avg Processes Blocked on Priority
+-----------------+
GPRAVPRQ contains the average number of processes waiting on
for their priority to become high enough to get the CPU
during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           GPRPRITM - Time Blocked on Priority

ELEMENT DERIVATION: GPRAVPRQ is computed using the following
algorithm:

  GPRAVPRQ = DURATION;
  IF DURATION
    THEN GPRAVPRQ = GPRPRITM/DURATION;

RANGE OF VALUE:  GPRAVPRQ takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVPRQ contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R A V R T M |  Avg Run Time of Completing Processes
+-----------------+
GPRAVRTM contains the average number of seconds of run time
accumulated by processes that completed during the
measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRRUNTM - Total Run Time of Completing Processes
           GPRCPRCT - Total Processes Completed

ELEMENT DERIVATION: GPRAVRTM is computed using the following
algorithm:

  GPRAVRTM = GPRCPRCT;
  IF GPRCPRCT
    THEN GPRAVRTM = GPRRUNTM/GPRCPRCT;

RANGE OF VALUE:  GPRAVRTM takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVRTM contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R A V S L Q |  Avg Processes in Sleep State
+-----------------+
GPRAVSLQ contains the average number of processes pausing or
sleeping as a result of their code calling sleep, pause, or
alarm instructions during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRSLQTM - Sleep Queue Wait Time
```

```
            DURATION - Recording Interval Time

ELEMENT DERIVATION: GPRAVSLQ is computed using the following
algorithm:

  GPRAVSLQ = DURATION;
  IF DURATION
    THEN GPRAVSLQ = GPRSLQTM/DURATION;

RANGE OF VALUE:  GPRAVSLQ takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVSLQ contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  On most systems, a significant percentage of
all processes will be in the sleep queue.  This is a normal
wait state for system daemons that periodically wake up and
poll status.  It does not indicate a resource bottleneck.


+-----------------+
| G P R A V S U S |  Avg Processes Suspended
+-----------------+
GPRAVSUS contains the average number of processes suspended
during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRSUSCT - Sum of Avg Processes Suspended
           AUMSAMPL - Number of Data Samples

ELEMENT DERIVATION: GPRAVSUS is computed using the following
algorithm:

  GPRAVSUS = AUMSAMPL;
  IF AUMSAMPL
    THEN GPRAVSUS = GPRSUSCT/AUMSAMPL;

RANGE OF VALUE:  GPRAVSUS takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVSUS contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| G P R A V S W A |  Avg Processes Waiting Swap In
+-----------------+
GPRAVSWA contains the average number of processes waiting to
be swapped-in during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRSWACT - Sum of Avg Processes Waiting swapin
           AUMSAMPL - Number of Data Samples

ELEMENT DERIVATION: GPRAVSWA is computed using the following
algorithm:

  GPRAVSWA = AUMSAMPL;
  IF AUMSAMPL
    THEN GPRAVSWA = GPRSWACT/AUMSAMPL;

RANGE OF VALUE:  GPRAVSWA takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRAVSWA contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R A V U S R |  Avg Number of Login Users
+-----------------+
GPRAVUSR contains the average number of users logged on to
the system during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  AUMSAMPL - Number of Data Samples
           GPRUSRCT - Sum of Avg Login Users

ELEMENT DERIVATION:  GPRAVUSR is computed using the following
algorithm:

  GPRAVUSR = AUMSAMPL;
  IF AUMSAMPL
    THEN GPRAVUSR = GPRUSRCT/AUMSAMPL;

RANGE OF VALUE:  GPRAVUSR takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.
```

SUMMARIZATION PROCESS: GPRAVUSR contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: Fluctuations in this average indicate that a user or application might make more efficient use of resources by maintaining its connection rather than logging onto and out of the system.

G P R B I O C T | Processes Blocked on Local Disks

GPRBIOCT is the accumulation of the average number of processes blocked on local resources.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: GPRAVBIQ- Avg Processes Blocked on Local Disks

ELEMENT DERIVATION: At the DETAIL time-span, GPRBIOCT is computed using MeasureWare Global metrics in the following calculation:

```
GBL_BLOCKED_IO_QUEUE*GBL_PROC_SAMPLE
```

RANGE OF VALUE: GPRBIOCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRBIOCT contains the sum of all values of GPRBIOCT processed in the level of summarization.

USAGE NOTES: None

G P R C A C H E | Size of Buffer Cache Configured

GPRCACHE contains the number of kilobytes of memory space configured for the buffer cache.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPRCACHE is set to the value of the MeasureWare Configuration metric TBL_BUFFER_CACHE_AVAIL.

RANGE OF VALUE: GPRCACHE takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRCACHE contains the sum of all values of GPRCACHE processed in the level of summarization.

USAGE NOTES: None

G P R C H U S E | The Size of Currently Used Buffers

GPRCHUSE contains the size, in kilobytes, of the sum of currently used buffers.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPRCHUSE is set to the value of the MeasureWare Global metric TBL_BUFFER_CACHE_USED.

RANGE OF VALUE: GPRCHUSE ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRCHUSE contains the sum of all values of GPRCHUSE processed in the level of summarization.

USAGE NOTES: GPRCHUSE is updated every 30 seconds or at the sampling interval.

G P R C P Q T M | CPU Queue Wait Time

GPRCPQTM contains the time spent in a CPU queue during the measurement interval. This element is the sum of the time a process was blocked on priority and the time a process actually ran in a processor.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: DURATION - Recording Interval Time
GPRAVCPQ - Avg Processes in Run Queue

ELEMENT DERIVATION: At the DETAIL time-span, GPRCPQTM is computed using MeasureWare Global metrics in the following calculation:

```
GBL_RUN_QUEUE*INTERVAL
```

RANGE OF VALUE: GPRCPQTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRCPQTM contains the sum of all values of GPRCPQTM processed in the level of summarization.

USAGE NOTES: None

G P R C P R C T | Total Processes Completed

GPRCPRCT contains the number of processes completed during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPRCPRCT is set to the value of the MeasureWare Global metric GBL_COMPLETED_PROC.

RANGE OF VALUE: GPRCPRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRCPRCT contains the sum of all values of GPRCPRCT processed in the level of summarization.

USAGE NOTES: None

G P R C P U C P | Total Available CPU Time

GPRCPUCP contains the amount of CPU time available during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: DURATION - Recording Interval Time
MCFNCPUS - Number of Processors Configured

ELEMENT DERIVATION: At the DETAIL time-span, GPRCPUCP is computed using MeasureWare metrics in the following calculation:

INTERVAL*GBL_NUM_CPU

where INTERVAL is from Global data and GBL_NUM_CPU is from Configuration data.

RANGE OF VALUE: GPRCPUCP is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRCPUCP contains the sum of all values of GPRCPUCP processed in the level of summarization.

USAGE NOTES: None

G P R C P U T M | Total CPU Busy

GPRCPUTM contains the total amount of CPU busy time during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: GPRSYSTM - CPU Busy in System Mode
GPRUSRTM - CPU Busy in User Mode
GPRNICTM - CPU Busy at Nice Priorities
GPRRTMTM - CPU Busy at Real Time Priorities
GPRCSWTM - CPU Busy for Context Switching
GPRINTTM - CPU Busy for Interrupt Handling

ELEMENT DERIVATION: GPRCPUTM is set to the value of the MeasureWare Global metric GBL_CPU_TOTAL_TIME.

RANGE OF VALUE: GPRCPUTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRCPUTM contains the sum of all values of GPRCPUTM processed in the level of summarization.

USAGE NOTES: None

G P R C S W T M | CPU Busy for Context Switching

GPRCSWTM contains the number of seconds that the CPU was used by the dispatcher during this measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: At the DETAIL time-span, GPRCSWTM is computed using MeasureWare metrics in the following calculation:

```
  GBL_CPU_CSWITCH_TIME*GBL_NUM_CPU
```

where GBL_CPU_CSWITCH_TIME is from Global data and GBL_NUM_CPU is from Configuration data.

RANGE OF VALUE: GPRCSWTM is a SAS time value ranging from 0 to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: GPRCSWTM contains the sum of all values of GPRCSWTM processed in the level of summarization.

USAGE NOTES: None

G P R D B K C T | Block I/Os

GPRDBKCT contains the total number of block I/Os during the measurement interval. These are physical IOs generated by file system access and do not include virtual memory IO, or IOs relating to raw disk access. These are IOs for inode and superblock updates that are handled through the buffer cache.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDBRCT - Block Reads
GPRDBWCT - Block Writes

ELEMENT DERIVATION: GPRDBKCT is set to the value of the MeasureWare Application metric GBL_DISK_BLOCK_IO.

RANGE OF VALUE: GPRDBKCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDBKCT contains the sum of all values of GPRDBKCT processed in the level of summarization.

USAGE NOTES: When a file is accessed, it is memory mapped by the operating system. Accesses generate virtual memory IOs. Reading a file generates block IOs as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM IOs) and posting updated inode information to disk (block IOs).

G P R D B R C T | Block Reads

GPRDBRCT contains the number of block reads during the measurement interval. These are physical I/Os generated by file system access and do not include virtual memory I/O or I/Os relating to raw disk access, on HP operating systems.

On Sun operating systems, these are I/Os for inode and superblock updates that are handled through the buffer cache.

On IBM operating systems, these are I/Os that do not include the read of the inode and file data.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDBWCT - Block Writes

ELEMENT DERIVATION: GPRDBRCT is set to the value of the MeasureWare Application metric GLB_DISK_BLOCK_READ.

RANGE OF VALUE: GPRDBRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDBRCT contains the sum of all values of GPRDBRCT processed in the level of summarization.

USAGE NOTES: When a file is accessed, it is memory mGPRed by the operating system. Accesses generate virtual memory I/Os. Reading a file generates block I/Os as the file's inode information is cached. File writes are a combination of posting to memory mGPRed allocations (VM I/Os) and posting updated inode information to disk (block I/Os).

G P R D B W C T | Block Writes

GPRDBWCT contains the number of block writes during the measurement interval. These are physical I/Os generated by file system access and do not include virtual memory I/O or I/Os relating to raw disk access, on HP operating systems.

These are I/Os for inode and superblock updates that are handled through the buffer cache on Sun operating systems.

For IBM operating systems, the block writes do not include the write of the inode and the file system data write.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDBRCT - Block Reads

ELEMENT DERIVATION: GPRDBWCT is set to the value of the MeasureWare Application metric GBL_DISK_BLOCK_WRITE.

RANGE OF VALUE: GPRDBWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDBWCT contains the sum of all values of GPRDBWCT processed in the level of summarization.

USAGE NOTES: When a file is accessed, it is memory mapped by the operating system. Accesses generate virtual memory I/Os. Reading a file generates block I/Os as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM I/Os) and posting updated inode information to disk (block I/Os).

G P R D F R C T | Disk Reads by File System

GPRDFRCT contains the number of disk reads by the file system during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDFWCT - Disk Writes by File System

ELEMENT DERIVATION: GPRDFRCT is set to the value of the MeasureWare Global metric GBL_DISK_FS_READ.

RANGE OF VALUE: GPRDFRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDFRCT contains the sum of all values of GPRDFRCT processed in the level of summarization.

USAGE NOTES: None

G P R D F W C T | Disk Writes by File System

GPRDFWCT contains the number of disk writes by the file system during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDFRCT - Disk Reads by File System

ELEMENT DERIVATION: GPRDFWCT is set to the value of the MeasureWare Global metric GBL_DISK_FS_WRITE.

RANGE OF VALUE: GPRDFWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDFWCT contains the sum of all values of GPRDFWCT processed in the level of summarization.

USAGE NOTES: None

G P R D K L R T | Logical Disk I/O Rate

GPRDKLRT contains the rate of logical disk I/Os during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time
GPRLDKCT - Logical Disk I/Os

ELEMENT DERIVATION: GPRDKLRT is computed using the following algorithm:

```
GPRDKLRT = DURATION;
IF DURATION
  THEN GPRDKLRT = GPRLDKCT/DURATION;
```

RANGE OF VALUE: GPRDKLRT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDKLRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: Logical I/Os may not always result in physical disk access since the data may already reside in the buffer cache.

G P R D K Q T M | Disk I/O Queue Wait Time

GPRDKQTM contains the time spent waiting in a disk I/O queue during the measurement interval. It includes processes blocked on disk, inode, CDFS, and general I/O waits.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: DURATION - Recording Interval Time
GPRAVDKQ - Avg Processes Waiting for Disk I/Os

ELEMENT DERIVATION: At the DETAIL time-span, GPRDKQTM is computed using MeasureWare Global metrics in the following calculation:

```
GBL_DISK_SUBSYSTEM_QUEUE*INTERVAL
```

RANGE OF VALUE: GPRDKQTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDKQTM contains the sum of all values of GPRDKQTM processed in the level of summarization.

USAGE NOTES: GPRDKQTM measures disk blocked time as an average of total execution time for all processes. While it is not a direct measure of disk usage, it does indicate when disk contention occurs among active processes.

G P R D M M C T | Virtual Memory Management I/Os

GPRDMMCT contains the number of I/Os completed for virtual memory management tasks during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDMMRT - Virtual Memory Management I/O Rate
GPRDMRCT - Disk Reads for Virt Memory Management
GPRDMWCT - Disk Writes for Virt Memory Management

ELEMENT DERIVATION: At the DETAIL timespan, GPRDMMCT is computed using MeasureWare Global metrics in the following calculation:

```
GBL_DISK_VM_READ+GBL_DISK_VM_WRITE
```

RANGE OF VALUE: GPRDMMCT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDMMCT contains the sum of all values of GPRDMMCT processed in the level of summarization.

USAGE NOTES: For the HP-UX platform, I/Os to the user file are not included in this metric unless they were done via the mmap(2) system call.

G P R D M M R T | Virtual Memory Management I/O Rate

GPRDMMRT contains the rate at which I/Os completed for virtual memory management tasks during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is integer.

SEE ALSO: DURATION - Recording Interval Time
GPRDMMCT - Virtual Memory Management I/Os

ELEMENT DERIVATION: GPRDMMRT is computed using the following algorithm:

```
GPRDMMRT = DURATION;
IF DURATION
  THEN GPRDMMRT = GPRDMMCT/DURATION;
```

RANGE OF VALUE: GPRDMMRT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDMMRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

G P R D M R C T | Disk Reads for Virtual Memory Management

GPRDMRCT contains the number of disk reads performed to manage virtual memory during the measurement interval. This includes page-ins (a process references code or data that is not in memory) and swap-ins (the system resumes execution of a process that was swapped out).

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRAVMMQ - Avg Processes Waiting for Memory
GPRDMMCT - Virtual Memory Management I/Os
GPRDMWCT - Disk Writes for Virt Memory Management
GPRPCMEM - Percent of Main Memory Used

ELEMENT DERIVATION: GPRDMRCT is set to the value of the

MeasureWare Global metric GBL_DISK_VM_READ.

RANGE OF VALUE: GPRDMRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDMRCT contains the sum of all values of GPRDMRCT processed in the level of summarization.

USAGE NOTES: Virtual memory management tasks include paging and swapping. If GPRDMRCT's value is high and GPRPCMEM's value nears 100, it indicates the system under memory pressure.

G P R D M R R T | Virt Memory Management Disk Read Rate

GPRDMRRT contains the rate at which disks are read for virtual memory management. This element measures the disk read rate for paging and swapping activity, including page-ins (a process references code or data that is not in memory) and swap-ins (the system resumes execution of a process that was swapped out).

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time
GPRDMRCT - Disk Reads for Virtual Memory Mgmt

ELEMENT DERIVATION: GPRDMRRT is computed using the following algorithm:

```
GPRDMRRT = DURATION;
IF DURATION
  THEN GPRDMRRT = GPRDMRCT/DURATION;
```

RANGE OF VALUE: GPRDMRRT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDMRRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: GPRDMRRT measures how well a system's physical memory is being used. If GPRDMRRT's value is high and GPRPCMEM's value nears 100, it indicates the system is under memory pressure.

G P R D M W C T | Disk Writes for Virtual Memory Mgmt

GPRDMWCT contains the number of disk writes performed to manage virtual memory during the measurement interval. It measures the disk writes for paging and swapping activity, including page-outs (the system moves an unused page out because of memory pressure) and swap-outs (the system moves all pages attached to a process from physical memory to the swap area).

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRAVMMQ - Avg Processes Waiting for Memory
GPRDMMCT - Virtual Memory Management I/Os
GPRDMRCT - Disk Reads for Virt Memory Management
GPRPCMEM - Percent of Main Memory Used

ELEMENT DERIVATION: GPRDMWCT is set to the value of the MeasureWare Global metric GBL_DISK_VM_WRITE.

RANGE OF VALUE: GPRDMWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDMWCT contains the sum of all values of GPRDMWCT processed in the level of summarization.

USAGE NOTES: Virtual memory management tasks include paging and swapping. If GPRDMWCT's value is high and GPRPCMEM's value nears 100, it indicates the system is under memory pressure. Additionally, writes to user file data aren't included in GPRDMWCT unless they were done via the mmap(2) system call.

G P R D M W R T | Virt Memory Management Disk Write Rate

GPRDMWRT contains the rate at which disks are written for virtual memory management. It measures the disk write rate for paging and swapping activity, including page-outs (the system moves an unused page out because of memory pressure) and swap-outs (the system moves all pages attached to a process from physical memory to the swap area).

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time
GPRDMMRT - Virtual Memory Management I/O Rate
GPRDMWCT - Disk Writes for Virtual Memory Mgmt

ELEMENT DERIVATION: GPRDMWRT is computed using the following algorithm:

```
  GPRDMWRT = DURATION;
  IF DURATION
    THEN GPRDMWRT = GPRDMWCT/DURATION;
```

RANGE OF VALUE: GPRDMWRT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDMWRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: GPRDMWRT measures how well a system's physical memory is being used. If GPRDMWRT's value is high and GPRPCMEM's value nears 100, it indicates the system is under memory pressure.

G P R D R R C T — Raw Reads

GPRDRRCT contains the number of direct-mode or raw-mode physical disk reads issued by user processes during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDRWCT - Raw Writes

ELEMENT DERIVATION: GPRDRRCT is set to the value of the MeasureWare Global metric GBL_DISK_RAW_READ.

RANGE OF VALUE: GPRDRRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDRRCT contains the sum of all values of GPRDRRCT processed in the level of summarization.

USAGE NOTES: A raw read bypasses the file system's buffering. Each raw read results in exactly one I/O.

G P R D R W C T — Raw Writes

GPRDRWCT contains the number of direct-mode or raw-mode disk writes issued by user processes during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDRRCT - Raw Reads

ELEMENT DERIVATION: GPRDRWCT is set to the value of the MeasureWare Global metric GBL_DISK_RAW_WRITE.

RANGE OF VALUE: GPRDRWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDRWCT contains the sum of all values of GPRDRWCT processed in the level of summarization.

USAGE NOTES: A raw write bypasses the file system's buffering. Each raw write results in exactly one I/O.

G P R D S K I O — Physical Disk I/Os

GPRDSKIO contains the number of physical disk I/Os during the measurement interval. Physical writes and reads to disk (including virtual memory and raw I/O) are included.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRPRDCT - Physical Disk Reads
GPRPWTCT - Physical Disk Writes

ELEMENT DERIVATION: GPRDSKIO is set to the value of the MeasureWare Global metric GBL_DISK_PHYS_IO.

RANGE OF VALUE: GPRDSKIO takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRDSKIO contains the sum of all values of GPRDSKIO processed in the level of summarization.

USAGE NOTES: On a SunOS platform, if a CD drive is not powered on or inserted into the CD drive at boottime; the system will not provide performance data for that device.

```
+-----------------+
| G P R D S K R T |  Physical Disk I/O Rate
+-----------------+
GPRDSKRT contains the rate of physical disk I/Os per second.
It measures the average rate per second of physical transfers
to and from all disk devices on the system.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           GPRDSKIO - Physical Disk I/Os

ELEMENT DERIVATION:  GPRDSKRT is computed using the following
algorithm:

  GPRDSKRT = DURATION;
  IF DURATION
    THEN GPRDSKRT = GPRDSKIO/DURATION;

RANGE OF VALUE:  GPRDSKRT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRDSKRT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  High GPRDSKRT values are possible on systems
with many disk drives, and on systems with fast disk drives.
Unusually large values for GPRDSKRT may indicate processes
that are inefficiently accessing files.


+-----------------+
| G P R D S R C T |  Disk Reads for File System Management
+-----------------+
GPRDSRCT contains the number of physical disk reads that took
place for file system management during the measurement
interval.  This is system overhead caused by read access to
inodes, cylinder groups, and superblocks.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRDFRCT - Disk Reads by File System
           GPRDMRCT - Disk Reads for Virt Memory Management
           GPRDRRCT - Raw Reads
           GPRDSWCT - Disk Writes for File System Management
           GPRPRDCT - Physical Disk Reads

ELEMENT DERIVATION:  GPRDSRCT is set to the value of the
MeasureWare Global metric GBL_DISK_SYSTEM_READ.

RANGE OF VALUE:  GPRDSRCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRDSRCT contains the sum of all
values of GPRDSRCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R D S W C T |  Disk Writes for System Use
+-----------------+
GPRDSWCT contains the number of times to which physical disks
were written for the system's use during the measurement
interval.  This is system overhead caused by write access to
inodes, cylinder groups, and superblocks.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRDFWCT - Disk Writes by File System
           GPRDMWCT - Disk Writes for Virt Memory Management
           GPRDRWCT - Raw Writes
           GPRDSRCT - Disk Reads for File System Management
           GPRPWTCT - Physical Disk Writes

ELEMENT DERIVATION:  GPRDSWCT is set to the value of the
MeasureWare Global metric GBL_DISK_SYSTEM_WRITE.

RANGE OF VALUE:  GPRDSWCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRDSWCT contains the sum of all
values of GPRDSWCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R F S D C T |  File System I/Os
+-----------------+
GPRFSDCT contains the number of local file system disk reads
and writes during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRFSMCT - File System Management I/Os

ELEMENT DERIVATION: GPRFSDCT is set to the value of the
MeasureWare Global metric GBL_DISK_FS_IO.
```

RANGE OF VALUE: GPRFSDCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRFSDCT contains the sum of all values of GPRFSDCT processed in the level of summarization.

USAGE NOTES: None

G P R F S M C T | File System Management I/Os

GPRFSMCT contains the number of local physical disk I/Os generated for file system management during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRFSDCT - File System I/Os

ELEMENT DERIVATION: GPRFSMCT is set to the value of the MeasureWare Global metric GBL_DISK_SYSTEM_IO.

RANGE OF VALUE: GPRFSMCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRFSMCT contains the sum of all values of GPRFSMCT processed in the level of summarization.

USAGE NOTES: None

G P R I N O D U | Inode Entries in Use

GPRINODU contains the number of "non-free" inodes currently in use on the HP-UX MeasureWare platform. GPRINODU contains the number of inode cache entries currently being used on the SOL MeasureWare platform.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPRINODU is set to the value of the MeasureWare Global metric TBL_INODE_CACHE_USED.

RANGE OF VALUE: GPRINODU ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRINODU contains the sum of all values of GPRINODU processed in the level of summarization.

USAGE NOTES: None

G P R I N T T M | CPU Busy for Interrupt Handling

GPRINTTM contains the number of seconds used by the CPU to process interrupts on behalf of user processes and idle CPU activity (clock handling, mux polling, etc.) during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: GPRPCINT - Pct CPU Busy for Interrupt Handling

ELEMENT DERIVATION: GPRINTTM is set to the value of the MeasureWare Global metric GBL_CPU_INTERRUPT_TIME.

RANGE OF VALUE: GPRINTTM is a SAS time value ranging from 0 to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: GPRINTTM contains the sum of all values of GPRINTTM processed in the level of summarization.

USAGE NOTES: None

G P R I P Q T M | IPC Queue Wait Time

GPRIPQTM contains the time spent waiting in an IPC queue during the measurement interval. The element reports the time that processes are blocked waiting for completion of semaphore, message queue, pipe, socket, or other IPC operations.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: DURATION - Recording Interval Time
GPRAVIPQ - Avg Processes Waiting for IPC Resources

ELEMENT DERIVATION: At the DETAIL time-span, GPRIPQTM is computed using MeasureWare Global metrics in the following calculation:

```
  GBL_IPC_SUBSYSTEM_QUEUE*INTERVAL

RANGE OF VALUE:  GPRIPQTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRIPQTM contains the sum of all
values of GPRIPQTM processed in the level of summarization.

USAGE NOTES:  None
```

G P R L C O L S — LAN Collisions

GPRLCOLS contains the number of LAN collisions that occured during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: DURATION - Recording Interval Time
GPRAVLCL - Rate of LAN Collisions per Minute

ELEMENT DERIVATION: At the DETAIL time-span, GPRLCOLS is computed using MeasureWare Global metrics in the following calculation:

```
  ROUND(GBL_NET_COLLISION_1_MIN_RATE*INTERVAL/60,1)
```

RANGE OF VALUE: GPRLCOLS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRLCOLS contains the sum of all values of GPRLCOLS processed in the level of summarization.

USAGE NOTES: Collisions occur on any busy LAN, but abnormal collision rates could indicate that too many packets are being generated due to a hardware or software problem.

G P R L D K C T — Logical Disk I/Os

GPRLDKCT contains the number of logical I/Os to all disk devices during the measurement interval. It measures both read and write system calls.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRLRDCT - Logical Disk Reads
GPRLWTCT - Logical Disk Writes

ELEMENT DERIVATION: GPRLDKCT is set to the value of the MeasureWare Global metric GBL_DISK_LOGL_IO.

RANGE OF VALUE: GPRLDKCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRLDKCT contains the sum of all values of GPRLDKCT processed in the level of summarization.

USAGE NOTES: Logical I/Os may not always result in physical disk access since the data may already reside in the buffer cache.

G P R L E R R S — LAN Packet Errors

GPRLERRS contains the number of LAN packet errors that occurred during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: DURATION - Recording Interval Time
GPRAVLER - Rate of LAN Errors per Minute

ELEMENT DERIVATION: At the DETAIL time-span, GPRLERRS is computed using MeasureWare Global metrics in the following calculation:

```
  ROUND(GBL_NET_ERROR_1_MIN_RATE*INTERVAL/60,1)
```

RANGE OF VALUE: GPRLERRS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRLERRS contains the sum of all values of GPRLERRS processed in the level of summarization.

USAGE NOTES: A large LAN packet error rate could indicate a hardware or software problem.

G P R L O S T B — Trace Buffers Lost

GPRLOSTB contains the number of trace buffers that are lost by the measurement processing daemon.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPRLOSTB is set to the value of the MeasureWare Global metric GBL_LOST_MI_TRACE_BUFFERS.

RANGE OF VALUE: GPRLOSTB ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRLOSTB contains the sum of all values of GPRLOSTB processed in the level of summarization.

USAGE NOTES: If the value of GPRLOSTB is greater than zero, this signifies that the measurement subsystem is not keeping up with system events on the HP-UX MeasureWare platform or ARM API calls on the SOL MeasureWare platform that generates traces.

G P R L P C K I | Successful LAN Packets Received

GPRLPCKI contains the number of LAN packets received through all network interfaces during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRLPCKO - Successful LAN Packets Sent

ELEMENT DERIVATION: GPRLPCKI is set to the value of the MeasureWare Global metric GBL_NET_IN_PACKET.

RANGE OF VALUE: GPRLPCKI takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRLPCKI contains the sum of all values of GPRLPCKI processed in the level of summarization.

USAGE NOTES: None

G P R L P C K O | Successful LAN Packets Sent

GPRLPCKO contains the number of LAN packets sent through all network interfaces during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRLPCKI - Successful LAN Packets Received

ELEMENT DERIVATION: GPRLPCKO is set to the value of the MeasureWare Global metric GBL_NET_OUT_PACKET.

RANGE OF VALUE: GPRLPCKO takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRLPCKO contains the sum of all values of GPRLPCKO processed in the level of summarization.

USAGE NOTES: None

G P R L R D C T | Logical Disk Reads

GPRLRDCT contains the number of logical disk reads during the measurement interval.

For the HP-UX platform, GPRLRDCT contains the number of logical reads during the measurement interval.

For the Sun platform, GPRLRDCT contains the number of logical block reads during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRLRDKB - Bytes Logically Read From Disk
GPRLWTCT - Logical Disk Writes

ELEMENT DERIVATION: GPRLRDCT is set to the value of the MeasureWare Global metric GBL_DISK_LOGL_READ.

RANGE OF VALUE: GPRLRDCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRLRDCT contains the sum of all values of GPRLRDCT processed in the level of summarization.

USAGE NOTES: Several possibilites may contribute to logical I/Os not corresponding with physical I/Os on the HP-UX platform. An example of this would be where a single logical write can cause a physical read to call the block to be updated from disk, and then cause a physical write to put it back on to the disk.

```
+-----------------+
| G P R L R D K B |  KBytes Logically Read From Disk
+-----------------+
GPRLRDKB contains the number of kilobytes of data logically
read from disk during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRLWTKB - KBytes Logically Written to Disk

ELEMENT DERIVATION:  GPRLRDKB is set to the value of the
MeasureWare Global metric GBL_DISK_LOGL_READ_BYTE.

RANGE OF VALUE:  GPRLRDKB takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRLRDKB contains the sum of all
values of GPRLRDKB processed in the level of summarization.

USAGE NOTES:  None



+-----------------+
| G P R L U S E D |  Number of File Locks in Use
+-----------------+
GPRLUSED contains the number of record or file locks that
are currently in use on SUN Solaris platforms.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  GPRLUSED is set to the value of the
MeasureWare Global metric TBL_FILE_LOCK_USED.

RANGE OF VALUE:  GPRLUSED ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRLUSED contains the sum of all
values of GPRLUSED processed in the level of summarization.

USAGE NOTES:  One file can have multiple locks.



+-----------------+
| G P R L W T C T |  Logical Disk Writes
+-----------------+
GPRLWTCT contains the number of logical disk writes during
the measurement interval.

For the HP-UX platform, GPRLWTCT contains the number of
logical writes during the measurement interval.

For the Sun platform, GPRLWTCT contains the number of logical
block writes during the measurement interval, where only
local disks are counted.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRLRDCT - Logical Disk Reads
           GPRLWTKB - KBytes Logically Written To Disk

ELEMENT DERIVATION:  GPRLWTCT is set to the value of the
MeasureWare Global metric GBL_DISK_LOGL_WRITE.

RANGE OF VALUE:  GPRLWTCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRLWTCT contains the sum of all
values of GPRLWTCT processed in the level of summarization.

USAGE NOTES:  Several possibilites may contribute to logical
I/Os not corresponding with physical I/Os on the HP-UX
platform.  An example of this would be where a single logical
write can cause a physical read to call the block to be
updated from disk, and then cause a physical write to put it
back on to the disk.



+-----------------+
| G P R L W T K B |  KBytes Logically Written To Disk
+-----------------+
GPRLWTKB contains the number of kilobytes of data logically
written to disk during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRLRDKB - KBytes Logical Read From Disk

ELEMENT DERIVATION:  GPRLWTKB is set to the value of the
MeasureWare Global metric GBL_DISK_LOGL_WRITE_BYTE.

RANGE OF VALUE:  GPRLWTKB takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRLWTKB contains the sum of all
values of GPRLWTKB processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| G P R M E M S W |  Amount of Swap Space Configured
+-----------------+
GPRMEMSW contains the number of kilobytes of disk space
configured for swapping.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  MCFMEMSW - Swap Space Configured

ELEMENT DERIVATION:  At the DETAIL time-span, GPRMEMSW is set
to the value of MCFMEMSW.

RANGE OF VALUE:  GPRMEMSW takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRMEMSW contains the sum of all
values of GPRMEMSW processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R M E M S Z |  Amount of Main Memory Configured
+-----------------+
GPRMEMSZ contains the amount of main memory (in kilobytes)
configured.  It equals the total capacity of all hardware
memory boards.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  MCFMEMSZ - Main Memory Configured

ELEMENT DERIVATION:  At the DETAIL time-span, GPRMEMSZ is set
to the value of MCFMEMSZ.

RANGE OF VALUE:  GPRMEMSZ takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRMEMSZ contains the sum of all
values of GPRMEMSZ processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R M E M U S |  Amount of Main Memory Available for User
+-----------------+
GPRMEMUS contains the number of kilobytes of main memory
available to end users.  This is the amount of main memory
not reserved by the system or kernel tables and resident
buffers.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRMEMSZ - Amount of Main Memory Configured
           MCFMEMUS - Main Memory Available for Users

ELEMENT DERIVATION:  At the DETAIL time-span, GPRMEMUS is set
to the value of MCFMEMUS.

RANGE OF VALUE:  GPRMEMUS takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRMEMUS contains the sum of all
values of GPRMEMUS processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R M E M U T |  Main Memory Used
+-----------------+
GPRMEMUT contains the amount of memory, in kilobytes, in use
during the measurement interval.  This includes system
memory, buffer cache, and user memory.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRMMEMS - Main Memory Used by System
           GPRMMEMU - Main Memory Used by Users

ELEMENT DERIVATION:  At the DETAIL time-span, GPRMEMUT is
computed using MeasureWare Global metrics in the following
calculation:

  GBL_MEM_PHYS*GBL_MEM_UTIL

RANGE OF VALUE:  GPRMEMUT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRMEMUT contains the sum of all
values of GPRMEMUT processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| G P R M M E M S |  Main Memory Used by System
+-----------------+
GPRMMEMS contains the amount of main memory (in kilobytes)
used by the system during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is real
number.

SEE ALSO:  GPRMMEMU - Main Memory Used by Users
           GPRMEMSZ - Amount of Main Memory Configured

ELEMENT DERIVATION:  At the DETAIL time-span, GPRMMEMS is
computed using the following calculation:

  (GBL_MEM_SYS_AND_CACHE_UTIL*GPRMEMSZ)/100

where GBL_MEM_SYS_AND_CACHE_UTIL is from MeasureWare Global
data and GPRMEMSZ is from the CA MICS database.

RANGE OF VALUE:  GPRMMEMS takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRMMEMS contains the sum of all
values of GPRMMEMS processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R M M E M U |  Main Memory Used by Users
+-----------------+
GPRMMEMU contains the amount of main memory (in kilobytes)
allocated to user code and data, including shared memory,
during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is real
number.

SEE ALSO:  GPRPCMEM - Percent of Main Memory Used
           GPRMEMSZ - Amount of Main Memory Configured

ELEMENT DERIVATION:  At the DETAIL time-span, GPRMMEMU is
computed using the metrics in the following calculation:

  (GBL_MEM_USER_UTIL*GPRMEMSZ)/100;

where GBL_MEM_USER_UTIL is from MeasureWare Global data and
GPRMEMSZ is from the CA MICS database.

RANGE OF VALUE:  GPRMMEMU takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRMMEMU contains the sum of all
values of GPRMMEMU processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R M M Q T M |  Memory Queue Wait Time
+-----------------+
GPRMMQTM contains the time spent waiting in a memory queue
during the measurement interval.  It reports the time
processes are blocked waiting for memory buffers or virtual
memory I/O (VM, cache, and MBUF).

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  DURATION - Recording Interval Time
           GPRAVMMQ - Avg Processes Waiting for Memory

ELEMENT DERIVATION:  At the DETAIL time-span, GPRMMQTM is
computed using the metrics in the following calculation:

  GBL_MEM_QUEUE*DURATION

where GBL_MEM_QUEUE is from MeasureWare Global metric data
and DURATION is in the CA MICS database.

RANGE OF VALUE:  GPRMMQTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRMMQTM contains the sum of all
values of GPRMMQTM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R M P B D U |  Percent Utilization of Busiest Disk
+-----------------+
For the busiest disk in the measurement interval, GPRMPBDU
contains the percentage of time that this busiest disk was
performing I/O transfers.

TYPE:  Maximum - Numeric, measurement unit is real number.
```

SEE ALSO: DURATION - Recording Interval Time
GPRMXDTM - Time Busiest Disk Was Busy

ELEMENT DERIVATION: At the DETAIL time-span, GPRMPBDU is computed using MeasureWare Global metrics in the following calculation:

```
100*GBL_DISK_TIME_PEAK/INTERVAL
```

RANGE OF VALUE: GPRMPBDU takes any numeric value ranging from 1 to 100.

SUMMARIZATION PROCESS: GPRMPBDU contains the maximum value of GPRMPBDU processed for the level of summarization.

USAGE NOTES: A peak disk utilization of more than 50 percent often indicates a disk bottleneck. In that situation, the performance of an application is degraded because it is unable to complete disk I/O quickly.

Slower disk devices will show a higher utilization with lower throughput than faster disk devices.

G P R M P D S U | Peak Disk Space Utilization

GPRMPDSU contains the percentage of disk space used by the file system that is the most full among all the file systems during the measurement interval. It does not include CDFS (CD-ROM File System) type file systems, which are always 100 percent full.

TYPE: Maximum - Numeric, measurement unit is real number.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPRMPDSU is set to the value of the MeasureWare Global metric GBL_FS_SPACE_UTIL_PEAK.

RANGE OF VALUE: GPRMPDSU takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRMPDSU contains the maximum value of GPRMPDSU processed for the level of summarization.

USAGE NOTES: This element's value can exceed 100% in the following circumstances:

- the root user may have allocated blocks in the reserved area
- dynamic file system swapping may have used the reserved area

Values close to or exceeding 100 mean that the file system is full. Significant changes in GPRMPDSU signify that large amounts of file space are being allocated or deallocated.

G P R M Q U S E | Message Queues in Use

GPRMQUSE contains the amount of configured message queues currently used during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: GPRMSGQU - Message Queues Available

ELEMENT DERIVATION: At the DETAIL timespan, GPRMQUSE is computed using the following calculation:

```
(TBL_MSG_TABLE_UTIL*GPRMSGQU)/100
```

where TBL_MSG_TABLE_UTIL is from MeasureWare Global data and GPRMSGQU is from the CA MICS database.

RANGE OF VALUE: GPRMQUSE takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRMQUSE contains the sum of all values of GPRMQUSE processed in the level of summarization.

USAGE NOTES: None

G P R M S G Q U | Message Queues Available

GPRMSGQU contains the number of the message queues that can be allocated on a system during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: MCFMSGQU - Message Queues Available

ELEMENT DERIVATION: At the DETAIL timespan, GPRMSGQU is set to the value of MCFMSGQU.

RANGE OF VALUE: GPRMSGQU ranges from a minimum of 0 to a

maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRMSGQU contains the last value of GPRMSGQU processed in the level of summarization.

USAGE NOTES: None

G P R M X C P U | Max Number of Processors Configured

GPRMXCPU contains the maximum number of processors configured during the measurement interval.

TYPE: Maximum - Numeric, measurement unit is integer.

SEE ALSO: MCFNCPUS - Number of Processors Configured

ELEMENT DERIVATION: At the DETAIL time-span, GPRMXCPU is set to the value of MCFNCPUS.

RANGE OF VALUE: GPRMXCPU takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRMXCPU contains the maximum value of GPRMXCPU processed for the level of summarization.

USAGE NOTES: None

G P R M X D S K | Max Number of Disk Drives Configured

GPRMXDSK contains the maximum number of disk drives configured during the measurement interval.

TYPE: Maximum - Numeric, measurement unit is integer.

SEE ALSO: MCFNDSKS - Number of Disk Drives Configured

ELEMENT DERIVATION: At the DETAIL time-span, GPRMXDSK is set to the value of MCFNDSKS.

RANGE OF VALUE: GPRMXDSK takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRMXDSK contains the maximum value of GPRMXDSK processed for the level of summarization.

USAGE NOTES: None

G P R M X D T M | Time Busiest Disk Was Busy

GPRMXDTM contains the number of seconds used by the busiest disk to perform I/O transfers during the measurement interval.

TYPE: Maximum - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time

ELEMENT DERIVATION: GPRMXDTM is set to the value of the MeasureWare Global metric GBL_DISK_TIME_PEAK.

RANGE OF VALUE: GPRMXDTM takes any numeric value ranging from 0 to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: GPRMXDTM contains the maximum value of GPRMXDTM processed for the level of summarization.

USAGE NOTES: Your system may be experiencing disk bottlenecks if the value of GPRMXDTM is about half the value of the measurement interval (DURATION).

G P R M X L A N | Max Number of LAN Interfaces Configured

GPRMXLAN contains the maximum number of LAN interfaces configured during the measurement interval.

TYPE: Maximum - Numeric, measurement unit is integer.

SEE ALSO: MCFNLANS - Number of LAN Interfaces Configured

ELEMENT DERIVATION: At the DETAIL time-span, GPRMXLAN is set to the value of MCFNLANS.

RANGE OF VALUE: GPRMXLAN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRMXLAN contains the maximum value of GPRMXLAN processed for the level of summarization.

USAGE NOTES: None

```
+-----------------+
| G P R N F S C T |  Network File System Requests
+-----------------+
GPRNFSCT contains the number of Network File System requests,
both inbound and outbound, that occurred during the
measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  GPRNFSCT is set to the value of the
MeasureWare Global metric GBL_NFS_CALL.

RANGE OF VALUE:  GPRNFSCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRNFSCT contains the sum of all
values of GPRNFSCT processed in the level of summarization.

USAGE NOTES:  In a Network File System, a local system may
have disks that a remote system is using through an
NFS-mount.  The local system, in turn, may have a remote disk
NFS-mounted to the local system.  Inbound requests reflect
activity on local disks caused by the remote system.
Outbound requests reflect activity by the local system on its
NFS-mounted remote disks.


+-----------------+
| G P R N F S R T |  Rate of Network File System Requests
+-----------------+
GPRNFSRT contains the number of network file system requests,
both inbound and outbound, per second during the measurement
interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           GPRNFSCT - Network File System Requests

ELEMENT DERIVATION:  GPRNFSRT is computed using the following
algorithm:

  GPRNFSRT = DURATION;
  IF DURATION
    THEN GPRNFSRT = GPRNFSCT/DURATION;

RANGE OF VALUE:  GPRNFSRT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRNFSRT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  In a Network File System, a local system may
have disks that a remote system is using through an
NFS-mount.  The local system, in turn, may have a remote disk
NFS-mounted to the local system.  Inbound requests reflect
activity on local disks caused by the remote system.
Outbound requests reflect activity by the local system on its
NFS-mounted remote disks.


+-----------------+
| G P R N I C T M |  CPU Busy at Nice Priorities
+-----------------+
GPRNICTM contains the number of seconds that the CPU was used
to run programs at a nice priority.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  GPRSYSTM - CPU Busy in System Mode

ELEMENT DERIVATION:  At the DETAIL time-span, GPRNICTM is
computed using MeasureWare metrics in the following
calculation:

  GBL_CPU_NICE_TIME*GBL_NUM_CPU

where GBL_CPU_NICE_TIME is from Global data and GBL_NUM_CPU
is from Configuration data.

RANGE OF VALUE:  GPRNICTM is a SAS time value ranging from 0
to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS:  GPRNICTM contains the sum of all
values of GPRNICTM processed in the level of summarization.

USAGE NOTES:  The Unix command 'nice' runs a command or
program at a non-default priority by specifying a nice value.
When the nice value increases, the process is given a smaller
share of available system resources.
```

```
+-----------------+
| G P R N O R T M |  CPU Busy at Normal Priorities
+-----------------+
GPRNORTM contains the number of seconds that the CPU was used
to run programs at normal priorities.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  GPRNICTM - CPU Busy at Nice Priorities
           GPRRTMTM - CPU Busy at Real Time Priorities

ELEMENT DERIVATION:  GPRNORTM is set to the value of the
MeasureWare Global metric GBL_CPU_NORMAL_TIME.

RANGE OF VALUE:  GPRNORTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRNORTM contains the sum of all
values of GPRNORTM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R N P R O C |  Value of NPROC Parameter
+-----------------+
GPRNPROC contains the size of the process table.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  MCFNPROC - Value of NPROC Parameter

ELEMENT DERIVATION:  At the DETAIL timespan, GPRNPROC is set
to the value of MCFNPROC.

RANGE OF VALUE:  GPRNPROC takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRNPROC contains the last value of
GPRNPROC processed in the level of summarization.

USAGE NOTES:  On the HP-UX platform, GPRNPROC is set by the
NPROC value during system generation.


+-----------------+
| G P R N T Q T M |  Network Queue Wait Time
+-----------------+
GPRNTQTM contains the time spent waiting in a network queue
for resources such as LAN, NFS (Network File System), DUX
(diskless), RFA (Remote File Access), and RPC (Remote
Procedure Call) to become available during the measurement
interval.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  DURATION - Recording Interval Time
           GPRAVNTQ - Avg Processes Waiting for Network

ELEMENT DERIVATION:  At the DETAIL time-span, GPRNTQTM is
computed using MeasureWare Global metrics in the following
calculation:

  GBL_NETWORK_SUBSYSTEM_QUEUE*INTERVAL

RANGE OF VALUE:  GPRNTQTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRNTQTM contains the sum of all
values of GPRNTQTM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R O F L O W |  Unmeasurable New Transactions
+-----------------+
GPROFLOW contains the number of new transactions that are
unmeasurable because the midaemon's Measurement Performance
Database is full.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  GPROFLOW is set to the value of the
MeasureWare Global metric GBL_TT_OVERFLOW_COUNT.

RANGE OF VALUE:  GPROFLOW ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPROFLOW contains the sum of all
values of GPROFLOW processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| G P R O U T Q U |  Outbound Queue Lengths
+-----------------+
GPROUTQU contains the sum of the outbound queue lengths for
all network interfaces within the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  GPROUTQU is set to the value of the
MeasureWare Global metric GBL_NET_OUTQUEUE.

RANGE OF VALUE:  GPROUTQU ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPROUTQU contains the sum of all
values of GPROUTQU processed in the level of summarization.

USAGE NOTES:  If the value of GPROUTQU is greater than zero
over a period of time, this may signify that the network is
experiencing bottlenecks.


+-----------------+
| G P R P A G R T |  Page Fault Rate
+-----------------+
GPRPAGRT contains the number of memory page faults (per
second) that occur during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           GPRPFLCT - Page Faults

ELEMENT DERIVATION:  GPRPAGRT is computed using the following
algorithm:

  GPRPAGRT = DURATION;
  IF DURATION
    THEN GPRPAGRT = GPRPFLCT/DURATION;

RANGE OF VALUE:  GPRPAGRT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRPAGRT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  A large paging rate may mean poor data
locality, excessive context switching, or insufficient
physical memory.


+-----------------+
| G P R P A U T M |  CPU Idle Due TO I/O Wait
+-----------------+
GPRPAUTM contains the time that the CPU was idle while
processes were waiting for physical IOs to complete during
the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  GPRCPUTM - Total CPU Busy

ELEMENT DERIVATION:  At the DETAIL time-span, GPRPAUTM is
computed using MeasureWare Global metrics in the following
calculation:

  GBL_CPU_WAIT_TIME*INTERVAL

RANGE OF VALUE:  GPRPAUTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRPAUTM contains the sum of all
values of GPRPAUTM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R P C A V M |  Percent of Virtual Memory in Active Use
+-----------------+
GPRPCAVM contains the percentage of virtual memory that was
used during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRMEMSW - Amount of Swap Space Configured
           GPRACTVM - Active Virtual Memory Size

ELEMENT DERIVATION:  GPRPCAVM is computed using the following
algorithm:

  GPRPCAVM = GPRMEMSW;
```

```
  IF GPRMEMSW
    THEN GPRPCAVM = 100*GPRACTVM/GPRMEMSW;
```

RANGE OF VALUE: GPRPCAVM takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: GPRPCAVM contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: GPRPCAVM values near 100 indicate that processes are efficiently using memory. GPRPCAVM values near 0 indicate that processes are reserving memory but not using it.

G P R P C C P U — Percent Total CPU Busy

GPRPCCPU contains the percentage of time that the CPU was busy during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRCPUCP - Total Available CPU Time
GPRCPUTM - Total CPU Busy

ELEMENT DERIVATION: GPRPCCPU is computed using the following algorithm:

```
  GPRPCCPU = GPRCPUCP;
  IF GPRCPUCP
    THEN GPRPCCPU = 100*GPRCPUTM/GPRCPUCP;
```

RANGE OF VALUE: GPRPCCPU can contain any value from 0 to a maximum limited only by practical bounds.

SUMMARIZATION PROCESS: GPRPCCPU contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

G P R P C C S W — Percent CPU Busy for Context Switching

GPRPCCSW contains the percentage of time that the CPU was used for context switching during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRCPUCP - Total Available CPU Time
GPRCSWTM - CPU Busy for Context Switching

ELEMENT DERIVATION: GPRPCCSW is computed using the following algorithm:

```
  GPRPCCSW = GPRCPUCP;
  IF GPRCPUCP
    THEN GPRPCCSW = 100*GPRCSWTM/GPRCPUCP;
```

RANGE OF VALUE: GPRPCCSW takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: GPRPCCSW contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

G P R P C H I T — Cache Hit Ratio

GPRPCHIT contains the cache hit ratio, which is the number of reads that were satisfied in the file system's buffer cache.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRRDHIT - Buffered Read Hits
GPRLRDCT - Logical Disk Reads

ELEMENT DERIVATION: GPRPCHIT is computed using the following algorithm:

```
  GPRPCHIT = GPRLRDCT;
  IF GPRLRDCT
    THEN GPRPCHIT = 100*GPRRDHIT/GPRLRDCT;
```

RANGE OF VALUE: GPRPCHIT takes any numeric value ranging from 0.0 to 100.0.

SUMMARIZATION PROCESS: GPRPCHIT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: A low cache hit ratio indicates that the buffer cache is not operating efficiently. Reasons include the size of the buffer cache (too small) and the organization of the data (large files, accessed sequentially).

A high hit ratio can result from an excessively large amount of memory configured for the file system buffer cache. This could lessen the physical memory available to user processes and result in excessive paging and swapping by the virtual memory management system. Thus, if you see a high hit ratio in the file system buffer cache, check for excessive activity in the virtual memory management system.

G P R P C I N T — Percent CPU Busy for Interrupt Handling

GPRPCINT contains the percentage of time that the CPU was busy handling interrupts on behalf of user processes and idle CPU activity (such as clock handling and mux polling) during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRCPUCP - Total Available CPU Time
GPRINTTM - CPU Busy for Interrupt Handling

ELEMENT DERIVATION: GPRPCINT is computed using the following algorithm:

```
GPRPCINT = GPRCPUCP
IF GPRCPUCP
  THEN GPRPCINT = 100*GPRINTTM/GPRCPUCP;
```

RANGE OF VALUE: GPRPCINT takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: GPRPCINT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

G P R P C M E M — Percent of Main Memory Used

GPRPCMEM contains the percentage of physical main memory allocated to user code and data during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRMEMSZ - Amount of Main Memory Configured
GPRMEMUT - Main Memory Used

ELEMENT DERIVATION: GPRPCMEM is computed using the following algorithm:

```
GPRPCMEM = GPRMEMSZ;
IF GPRMEMSZ
  THEN GPRPCMEM = 100*GPRMEMUT/GPRMEMSZ;
```

RANGE OF VALUE: GPRPCMEM takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: GPRPCMEM contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: If GPRPCMEM is near 100 and the value of GPRDMRRT or GPRDMWRT is high, it indicates the system is under memory pressure.

G P R P C M M S — Percent of Main Memory Used by System

GPRPCMMS contains the percentage of physical main memory used by the system during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRMEMSZ - Amount of Main Memory Configured
GPRMMEMS - Main Memory Used by System

ELEMENT DERIVATION: GPRPCMMS is computed using the following algorithm:

```
GPRPCMMS = GPRMEMSZ;
IF GPRMEMSZ
  THEN GPRPCMMS = 100*GPRMMEMS/GPRMEMSZ;
```

RANGE OF VALUE: GPRPCMMS takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: GPRPCMMS contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: If GPRPCMMS is high, it indicates that the majority of memory is not used for production work. Users may report poor response times.

G P R P C M M U — Percent of Main Memory Used by Users

GPRPCMMU contains the percentage of main memory used by users during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRMMEMU - Main Memory Used by Users
GPRMEMSZ - Amount of Main Memory Configured

ELEMENT DERIVATION: GPRPCMMU is computed using the following algorithm:

```
GPRPCMMU = GPRMEMSZ;
IF GPRMEMSZ
  THEN GPRPCMMU = 100*GPRMMEMU/GPRMEMSZ;
```

RANGE OF VALUE: GPRPCMMU takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: GPRPCMMU contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

G P R P C N I C — Percent CPU Busy at Nice Priorities

GPRPCNIC contains the percentage of time that the CPU was busy handling requests whose priority was NICE.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRCPUCP - Total Available CPU Time
GPRNICTM - CPU Busy at Nice Priorities

ELEMENT DERIVATION: GPRPCNIC is computed using the following algorithm:

```
GPRPCNIC = GPRCPUCP;
IF GPRCPUCP
  THEN GPRPCNIC = 100*GPRNICTM/GPRCPUCP;
```

RANGE OF VALUE: GPRPCNIC takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: GPRPCNIC contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

G P R P C N O R — Percent CPU Busy at Normal Priorities

GPRPCNOR contains the percentage of the CPU used to run programs at normal priorities during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: GPRNORTM - CPU Busy at Normal Priorities
GPRCPUCP - Total Available CPU Time

ELEMENT DERIVATION: GPRPCNOR is computed using the following algorithm:

```
GPRPCNOR = GPRCPUCP;
IF GPRCPUCP
  THEN GPRPCNOR = 100*GPRNORTM/GPRCPUCP;
```

RANGE OF VALUE: GPRPCNOR takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: GPRPCNOR contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

```
+-----------------+
| G P R P C R T M |  Percent CPU Busy at Real Time Priorities
+-----------------+
GPRPCRTM contains the percentage of time that the CPU was
busy satisfying requests whose priority fell in the real time
range.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRCPUCP - Total Available CPU Time
           GPRRTMTM - CPU Busy at Real Time Priorities

ELEMENT DERIVATION:  GPRPCRTM is computed using the following
algorithm:

  GPRPCRTM = GPRCPUCP;
  IF GPRCPUCP
    THEN GPRPCRTM = 100*GPRRTMTM/GPRCPUCP;

RANGE OF VALUE:  GPRPCRTM takes any numeric value ranging
from 0 to 100.

SUMMARIZATION PROCESS:  GPRPCRTM contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  Processes running at real time priorities do
not relinquish the CPU to processes running at lower
priorities.  If GPRPCRTM is large, your regular users may
perceive response time delays.


+-----------------+
| G P R P C S S C |  Percent CPU Busy for System Calls
+-----------------+
GPRPCSSC contains the percentage of the CPU used to satisfy
system calls during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRSSCTM - CPU Busy for System Calls
           GPRCPUCP - Total Available CPU Time

ELEMENT DERIVATION:  GPRPCSSC is computed using the following
algorithm:

  GPRPCSSC = GPRCPUCP;
  IF GPRCPUCP
    THEN GPRPCSSC = 100*GPRSSCTM/GPRCPUCP;

RANGE OF VALUE:  GPRPCSSC takes any numeric value ranges from
0 to 100.

SUMMARIZATION PROCESS:  GPRPCSSC contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R P C S W S |  Percent of Swap Space Used
+-----------------+
GPRPCSWS contains the percentage of swap space reserved for
running processes during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRMEMSW - Amount of Swap Space Configured
           GPRSWAPU - Swap Space Used

ELEMENT DERIVATION:  GPRPCSWS is computed using the following
algorithm:

  GPRPCSWS = GPRMEMSW;
  IF GPRMEMSW
    THEN GPRPCSWS = 100*GPRSWAPU/GPRMEMSW;

RANGE OF VALUE:  GPRPCSWS takes any numeric value ranging
from 0 to 100.

SUMMARIZATION PROCESS:  GPRPCSWS contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  When a process is created, the space need for
swap is reserved.  If there is not enough unreserved swap
space available as the value of GPRPCSWS nears 100, process
creation fails.  As the value of GPRPCSWS nears 0, consider
reducing the amount of swap space on your system to better
use disk space.
```

```
+-----------------+
| G P R P C S Y S |  Percent CPU Busy in System Mode
+-----------------+
GPRPCSYS contains the percentage of time that the system
operated in system mode.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRCPUCP - Total Available CPU Time
           GPRSYSTM - CPU Busy in System Mode

ELEMENT DERIVATION:  GPRPCSYS is computed using the following
algorithm:

  GPRPCSYS = GPRCPUCP;
  IF GPRCPUCP
    THEN GPRPCSYS = 100*GPRSYSTM/GPRCPUCP;

RANGE OF VALUE:  GPRPCSYS takes any numeric value ranging
from 0 to 100.

SUMMARIZATION PROCESS:  GPRPCSYS contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R P C U S R |  Percent CPU Busy in User Mode
+-----------------+
GPRPCUSR contains the amount of time the system spent
executing in user mode.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  GPRCPUCP - Total Available CPU Time
           GPRUSRTM - CPU Busy in User Mode

ELEMENT DERIVATION:  GPRPCUSR is computed using the following
algorithm:

  GPRPCUSR = GPRCPUCP;
  IF GPRCPUCP
    THEN GPRPCUSR = 100*GPRUSRTM/GPRCPUCP;

RANGE OF VALUE:  GPRPCUSR takes any numeric value ranging
from 0 to 100.

SUMMARIZATION PROCESS:  GPRPCUSR contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R P F L C T |  Page Faults
+-----------------+
GPRPFLCT contains the number of page faults recorded during
the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRPAGRT - Page Fault Rate

ELEMENT DERIVATION:  GPRPFLCT is set to the value of the
MeasureWare Global metric GBL_MEM_PAGE_REQUEST.

RANGE OF VALUE:  GPRPFLCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRPFLCT contains the sum of all
values of GPRPFLCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R P G O C T |  Number of Page Outs
+-----------------+
GPRPGOCT contains the number of page outs to disk during the
measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  GPRPGOCT is set to the value of the
MeasureWare Global metric GBL_MEM_PAGEOUT.

RANGE OF VALUE:  GPRPGOCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRPGOCT contains the sum of all
values of GPRPGOCT processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| G P R P K B C T |  KBytes Physically Transferred
+-----------------+
GPRPKBCT contains the number of kilobytes of data transferred
to and from local disks during the measurement interval.
This includes raw reads, reads issued by file systems, system
reads, and reads for virtual memory management.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  GPRPKBCT is set to the value of the
MeasureWare Global metric GBL_DISK_PHYS_BYTE.

RANGE OF VALUE:  GPRPKBCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRPKBCT contains the sum of all
values of GPRPKBCT processed in the level of summarization.

USAGE NOTES:  Spikes in this measurement can indicate a disk
bottleneck.

On a SunOS platform, if a CD drive is not powered on or
inserted into the CD drive at boottime, the system will not
provide performance data for that device.


+-----------------+
| G P R P R D C T |  Physical Disk Reads
+-----------------+
GPRPRDCT contains the number of physical disk reads during
the measurement interval.  This includes raw reads, reads
issued by file systems, system reads, and reads for virtual
memory management.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRPWTCT - Physical Disk Writes

ELEMENT DERIVATION:  GPRPRDCT is set to the value of the
MeasureWare Global metric GBL_DISK_PHYS_READ.

RANGE OF VALUE:  GPRPRDCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRPRDCT contains the sum of all
values of GPRPRDCT processed in the level of summarization.

USAGE NOTES:  On a SunOS platform, if a CD drive is not
powered on or inserted into the CD drive at boottime, the
system will not provide performance data for that device.

On an HP-UX platform there are multiple possibilities as to
why there may not be a direct connection between the number
of physical and logical I/Os.  An example of this would be
where small sequential logical records may be satisfied from
the buffer cache, resulting in fewer physical I/Os.


+-----------------+
| G P R P R I T M |  Time Blocked on Priority
+-----------------+
GPRPRITM contains the time that processes waited for their
priority to become high enough to get the CPU during the
measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  DURATION - Recording Interval Time
           GPRAVPRQ - Avg Processes Blocked on Priority

ELEMENT DERIVATION:  At the DETAIL time-span, GPRPRITM is
computed using MeasureWare Global metrics in the following
calculation:

  GBL_PRI_QUEUE*INTERVAL

RANGE OF VALUE:  GPRPRITM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRPRITM contains the sum of all
values of GPRPRITM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R P R O C T |  Sum of Avg Processes Alive
+-----------------+
GPRPROCT contains the sum of the average number of processes
alive that were sampled during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is real
number.

SEE ALSO:  AUMSAMPL - Number of Data Samples
```

GPRAVPRO - Avg Number of Processes Alive

ELEMENT DERIVATION: At the DETAIL time-span, GPRPROCT is computed using MeasureWare Global metrics in the following calculation:

```
GBL_ALIVE_PROC*GBL_PROC_SAMPLE
```

RANGE OF VALUE: GPRPROCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRPROCT contains the sum of all values of GPRPROCT processed in the level of summarization.

USAGE NOTES: GPRPROCT is used for computing GPRAVPRO only.

G P R P T B L U | Proc Table Entries Used

GPRPTBLU contains the amount of proc table entries currently used by the processes during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: GPRNPROC - Value of NPROC Parameter

ELEMENT DERIVATION: At the DETAIL timespan, GPRPTBLU is computed using the following calculation:

```
(TBL_PROC_TABLE_UTIL*GPRNPROC)/100
```

where TBL_PROC_TABLE_UTIL is from MeasureWare Global data and GPRNPROC is from the CA MICS database.

RANGE OF VALUE: GPRPTBLU takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRPTBLU contains the sum of all values of GPRPTBLU processed in the level of summarization.

USAGE NOTES: None

G P R P W T C T | Physical Disk Writes

GPRPWTCT contains the number of physical disk writes during the measurement interval. This includes raw writes, writes issued by file systems, system writes, and writes for virtual memory management.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRPRDCT - Physical Disk Reads

ELEMENT DERIVATION: GPRPWTCT is set to the value of the MeasureWare Global metric GBL_DISK_PHYS_WRITE.

RANGE OF VALUE: GPRPWTCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRPWTCT contains the sum of all values of GPRPWTCT processed in the level of summarization.

USAGE NOTES: On a SunOS platform, if a CD drive is not powered on or inserted into the CD drive at boottime, the system will not provide performance data for that device.

On an HP-UX platform there are multiple possibilities as to why there may not be a direct connection between the number of physical and logical I/Os. For example, small logical writes could end up entirely in the buffer cache and later generate fewer physical I/Os when written to disk due to the larger I/O size.

G P R R A W C T | Raw I/Os

GPRRAWCT contains the number of raw reads and writes to local disks during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRDRRCT - Raw Reads
GPRDRWCT - Raw Writes

ELEMENT DERIVATION: GPRRAWCT is set to the value of the MeasureWare Global metric GBL_DISK_RAW_IO.

RANGE OF VALUE: GPRRAWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRRAWCT contains the sum of all

values of GPRRAWCT processed in the level of summarization.

USAGE NOTES: None

G P R R D H I T | Buffered Read Hits

GPRRDHIT contains the number of requests satisfied by buffered storage during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRLRDCT - Logical Disk Reads
GPRPCHIT - Cache Hit Rate

ELEMENT DERIVATION: At the DETAIL time-span, GPRRDHIT is computed using MeasureWare Global metrics in the following calculation:

```
(GBL_MEM_CACHE_HIT_PCT*GBL_DISK_LOGL_READ)/100
```

RANGE OF VALUE: GPRRDHIT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRRDHIT contains the sum of all values of GPRRDHIT processed in the level of summarization.

USAGE NOTES: None

G P R R T M T M | CPU Busy at Real Time Priorities

GPRRTMTM contains the number of seconds that the CPU was used to run programs at real time priorities.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: GPRNICTM - CPU Busy at Nice Priorities

ELEMENT DERIVATION: At the DETAIL time-span, GPRRTMTM is computed using MeasureWare metrics in the following calculation:

```
GBL_CPU_REALTIME_TIME*GBL_NUM_CPU
```

where GBL_CPU_REALTIME_TIME is from Global data and GBL_NUM_CPU is from Configuration data.

RANGE OF VALUE: GPRRTMTM is a SAS time value ranging from 0 to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: GPRRTMTM contains the sum of all values of GPRRTMTM processed in the level of summarization.

USAGE NOTES: Processes running at real time priorities do not relinquish the CPU to processes running at lower priorities. If GPRRTMTM is large relative to the length of the measurement interval (data element DURATION), your regular users may perceive response time delays.

G P R R U N T M | Total Run Time of Completing Processes

GPRRUNTM contains the total number of seconds of run time accumulated by processes that completed during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: GPRAVRTM - Avg Run Time of Completing Processes

ELEMENT DERIVATION: At the DETAIL timespan, GPRRUNTM is computed using MeasureWare Global metrics in the following calculation:

```
GBL_PROC_RUN_TIME*GBL_COMPLETED_PROC
```

RANGE OF VALUE: GPRRUNTM is a SAS time value that ranges from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRRUNTM contains the sum of all values of GPRRUNTM processed in the level of summarization.

USAGE NOTES: None

G P R S C A L L | Total System Calls

GPRSCALL contains the total number of system calls received during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: DURATION - Recording Interval Time

GPRSCLRT - System Procedure Call Rate

ELEMENT DERIVATION: At the DETAIL time-span, GPRSCALL is computed using MeasureWare Global metrics in the following calculation:

```
GBL_SYSCALL_RATE*INTERVAL
```

RANGE OF VALUE: GPRSCALL takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSCALL contains the sum of all values of GPRSCALL processed in the level of summarization.

USAGE NOTES: High system call rates are normal on busy systems, especially with I/O-intensive applications. Abnormally high system call rates may indicate programs that are not using system calls efficiently.

G P R S C L R T | System Procedure Call Rate

GPRSCLRT contains the number of system procedure calls per second during the measurement interval. The system call count includes calls from all processes.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time
GPRDKLRT - Logical Disk I/O Rate
GPRSCALL - Total System Calls

ELEMENT DERIVATION: GPRSCLRT is computed using the following algorithm:

```
GPRSCLRT = DURATION;
IF DURATION
  THEN GPRSCLRT = GPRSCALL/DURATION;
```

RANGE OF VALUE: GPRSCLRT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSCLRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: System calls are requests to the operating system. High system call rates can be caused by inefficient use of system calls or I/O intensive processing.

Compare GPRSCLRT to GPRDKLRT to see if high system call rates correspond to high disk I/O rates.

G P R S E M I D | Semaphore Identifiers Available

GPRSEMID contains the number of semaphore identifiers that can be allocated on the system.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: MCFSEMID - Semaphore Identifiers Available

ELEMENT DERIVATION: At the DETAIL timespan, GPRSEMID is set to the value of MCFSEMID.

RANGE OF VALUE: GPRSEMID ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSEMID contains the last value of GPRSEMID processed in the level of summarization.

USAGE NOTES: None

G P R S E M S T | Semaphore Identifiers in Use

GPRSEMST contains the amount of configured semaphore identifiers currently used by the system during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRSEMID - Semaphore Identifiers Available

ELEMENT DERIVATION: At the DETAIL timespan, GPRSEMST is computed using the following calculation:

```
(TBL_SEM_TABLE_UTIL*GPRSEMID)/100
```

where TBL_SEM_TABLE_UTIL is from MeasureWare Global data and GPRSEMID is from the CA MICS database.

RANGE OF VALUE: GPRSEMST takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSEMST contains the sum of all

values of GPRSEMST processed in the level of summarization.

USAGE NOTES: None

G P R S H M S G | Table Memory Segments Available

GPRSHMSG contains the number of shared memory segments that can be allocated on a system during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: MCFSHMSG - Table Memory Segments Available

ELEMENT DERIVATION: GPRSHMSG is set to the value MCFSHMSG from the CA MICS database.

RANGE OF VALUE: GPRSHMSG ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSHMSG contains the last value of GPRSHMSG processed in the level of summarization.

USAGE NOTES: None

G P R S L Q T M | Sleep Queue Wait Time

GPRSLQTM contains the time spent waiting in a sleep queue during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: DURATION - Recording Interval Time
GPRAVSLQ - Avg Processes in Sleep State

ELEMENT DERIVATION: At the DETAIL time-span, GPRSLQTM is computed using MeasureWare Global metrics in the following calculation:

```
GBL_SLEEP_QUEUE*INTERVAL
```

RANGE OF VALUE: GPRSLQTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSLQTM contains the sum of all values of GPRSLQTM processed in the level of summarization.

USAGE NOTES: None

G P R S M E M S | Shared Memory Segments in Use

GPRSMEMS contains the amount of shared memory segments currently used by the system during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRSHMSG - Table Memory Segments Available

ELEMENT DERIVATION: At the DETAIL timespan, GPRSMEMS is computed using the following calculation:

```
(TBL_SHMEM_TABLE_UTIL*GPRSHMSG)/100
```

where GBL_SHMEM_TABLE_UTIL is from MeasureWare Global data and GPRSHMSG is from the CA MICS database.

RANGE OF VALUE: GPRSMEMS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSMEMS contains the sum of all values of GPRSMEMS processed in the level of summarization.

USAGE NOTES: None

G P R S P R C T | Total Processes Started

GPRSPRCT contains the number of processes started during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRCPRCT - Total Processes Completed

ELEMENT DERIVATION: GPRSPRCT is set to the value of the MeasureWare Global metric GBL_STARTED_PROC.

RANGE OF VALUE: GPRSPRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSPRCT contains the sum of all values of GPRSPRCT processed in the level of summarization.

USAGE NOTES: None

| G P R S S C T M | CPU Busy for System Calls

GPRSSCTM contains the number of seconds that the CPU was in system mode during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: GPRUSRTM - CPU Busy in User Mode
GPRNORTM - CPU Busy at Normal Priorities
GPRNICTM - CPU Busy at Nice Priorities
GPRRTMTM - CPU Busy at Real Time Priorities
GPRSYSTM - CPU Busy in System Mode
GPRCSWTM - CPU Busy for Context Switching
GPRINTTM - CPU Busy for Interrupt Handling
GPRCPUTM - Total CPU Busy

ELEMENT DERIVATION: GPRSSCTM is set to the value of the MeasureWare Global metric GBL_CPU_SYSCALL_TIME.

RANGE OF VALUE: GPRSSCTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSSCTM contains the sum of all values of GPRSSCTM processed in the level of summarization.

USAGE NOTES: None

| G P R S U P T M | Time of Last System Reboot

GPRSUPTM contains the time since the system was last rebooted.

TYPE: Retained - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPROFLOW is set to the value of the MeasureWare Global metric GBL_SYSTEM_UPTIME_HOURS.

RANGE OF VALUE: GPROFLOW is a SAS time value ranging from 0 to a maximum that is limited to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: GPRSUPTM contains the last value of GPRSUPTM processed in the level of summarization.

USAGE NOTES: None

| G P R S U S C T | Sum of Avg Suspended Processes

GPRSUSCT contains the sum of the average suspended processes during this measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: GPRAVSUS - Avg Processes Suspended
AUMSAMPL - Number of Data Samples

ELEMENT DERIVATION: At the DETAIL timespan, GPRSUSCT is computed using MeasureWare Global metrics in the following calculation:

GBL_SUSPENDED_PROCS*GBL_PROC_SAMPLE

RANGE OF VALUE: GPRSUSCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSUSCT contains the sum of all values of GPRSUSCT processed in the level of summarization.

USAGE NOTES: None

| G P R S W A C T | Sum of Avg Processes Awaiting Swap-In

GPRSWACT contains the sum of the average processes waiting to be swapped in during this measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: GPRAVSWA - Avg Processes Waiting Swap-In
AUMSAMPL - Number of Data Samples

ELEMENT DERIVATION: At the DETAIL timespan, GPRSWACT is computed using MeasureWare Global metrics in the following calculation:

```
  GBL_MEM_SWAP_QUEUE*GBL_PROC_SAMPLE

RANGE OF VALUE:  GPRSWACT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRSWACT contains the sum of all
values of GPRSWACT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R S W A P U |  Swap Space Used
+-----------------+
GPRSWAPU contains the amount of swap space in kilobytes used
during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is real
number.

SEE ALSO:  GPRMEMSW - Amount of Swap Space Configured
           GPRPCSWS - Percent of Swap Space Used

ELEMENT DERIVATION:  At the DETAIL time-span, GPRSWAPU is
computed using MeasureWare metrics in the following
calculation:

  (GBL_SWAP_SPACE_UTIL*GBL_SWAP_SPACE_AVAIL_KB)/100

where GBL_SWAP_SPACE_UTIL is from Global data and
GBL_SWAP_SPACE_AVAIL_KB is from Configuration data.

RANGE OF VALUE:  GPRSWAPU takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRSWAPU contains the sum of all
values of GPRSWAPU processed in the level of summarization.

USAGE NOTES:  When a process is created, the space need for
swap is reserved.  If there is not enough unreserved swap
space available as the value of GPRSWAPU nears the value of
GPRMEMSW, process creation fails.  As the value of GPRSWAPU
nears 0, consider reducing the amount of swap space on your
system to better use disk space.


+-----------------+
| G P R S W O C T |  Number of Swap Outs
+-----------------+
GPRSWOCT contains the number of process deactivations or swap
outs during this measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  GPRSWORT - Swap-Out Rate

ELEMENT DERIVATION:  At the DETAIL time-span, GPRSWOCT is set
to the value of the MeasureWare Global metrics in the
following calculation:

  GBL_MEM_SWAPOUT_RATE*INTERVAL

RANGE OF VALUE:  GPRSWOCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRSWOCT contains the sum of all
values of GPRSWOCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| G P R S W O R T |  Swap-Out Rate
+-----------------+
GPRSWORT contains the number of process deactivations or swap
outs per second during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is integer.

SEE ALSO:  DURATION - Recording Interval Time
           GPRSWOCT - Number of Swap Outs

ELEMENT DERIVATION:  GPRSWORT is computed using the following
algorithm:

  GPRSWORT = DURATION;
  IF DURATION
    THEN GPRSWORT = GPRSWOCT/DURATION;

RANGE OF VALUE:  GPRSWORT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  GPRSWORT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.
```

USAGE NOTES: None

| G P R S W P C T | Process Swaps

GPRSWPCT contains the number of swap-ins and swap-outs for all processes during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: GPRSWPCT is set to the value of the MeasureWare Global metric GBL_MEM_SWAP.

RANGE OF VALUE: GPRSWPCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSWPCT contains the sum of all values of GPRSWPCT processed in the level of summarization.

USAGE NOTES: If GPRSWPCT continually increases, your system may be experiencing memory constraints.

| G P R S W P R T | Swapping Rate

GPRSWPRT contains number of swap-ins and swap-outs per second for all processes during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is integer.

SEE ALSO: DURATION - Recording Interval Time
GPRSWPCT - Process Swaps, Activation/Deactivation

ELEMENT DERIVATION: GPRSWPRT is computed using the following algorithm:

```
GPRSWPRT = DURATION;
IF DURATION
  THEN GPRSWPRT = GPRSWPCT/DURATION;
```

RANGE OF VALUE: GPRSWPRT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRSWPRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

| G P R S Y S T M | CPU Busy in System Mode

GPRSYSTM contains the number of seconds that the CPU was used in system mode during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: GPRUSRTM - CPU Busy in User Mode

ELEMENT DERIVATION: At the DETAIL time-span, GPRSYSTM is computed using MeasureWare metrics in the following calculation:

```
GBL_CPU_SYS_MODE_TIME*GBL_NUM_CPU
```

where GBL_CPU_SYS_MODE_TIME is from Global data and GBL_NUM_CPU is from Configuration data.

RANGE OF VALUE: GPRSYSTM is a SAS time value ranging from 0 to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: GPRSYSTM contains the sum of all values of GPRSYSTM processed in the level of summarization.

USAGE NOTES: None

| G P R T B L E N | File Table Entries Used

GPRTBLEN contains the number of file table entries that are currently used by file descriptors during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: GPRTBLUK - Entries Used by Kernel

ELEMENT DERIVATION: At the DETAIL timespan, GPRTBLEN is computed using MeasureWare metrics in the following calculation:

```
(TBL_FILE_TABLE_UTIL*TBL_FILE_TABLE_AVAIL)/100
```

where both TBL_FILE_TABLE_UTIL is from Global data and TBL_FILE_TABLE_AVAIL is from Configuration data.

RANGE OF VALUE: GPRTBLEN ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRTBLEN contains the sum of all values of GPRTBLEN processed in the level of summarization.

USAGE NOTES: None

G P R T B L U K | Entries Used by Kernel

GPRTBLUK is the number of file table entries used by the kernel to manage open file descriptors.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: MCFTBLUK - Max Entries Used by Kernel

ELEMENT DERIVATION: At the DETAIL timespan, GPRTBLUK is set to the value of MCFTBLUK.

RANGE OF VALUE: GPRTBLUK takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRTBLUK contains the sum of all values of GPRTBLUK processed in the level of summarization.

USAGE NOTES: None

G P R U S R C T | Sum of Avg Login Users

GPRUSRCT is the accumulation of the average number of users logged on during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: AUMSAMPL - Number of Data Samples
GPRAVUSR - Avg Number of Login Users

ELEMENT DERIVATION: At the DETAIL time-span, GPRUSRCT is computed using MeasureWare Global metrics in the following calculation:

GBL_NUM_USER*GBL_PROC_SAMPLE

RANGE OF VALUE: GPRUSRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: GPRUSRCT contains the sum of all values of GPRUSRCT processed in the level of summarization.

USAGE NOTES: GPRUSRCT is used for computing the value of GPRAVUSR only.

G P R U S R T M | CPU Busy in User Mode

GPRUSRTM contains the number of seconds that the CPU was used in user mode during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: GPRSYSTM - CPU Busy in System Mode

ELEMENT DERIVATION: GPRUSRTM is set to the value of the MeasureWare Global metric GBL_CPU_USER_MODE_TIME.

RANGE OF VALUE: GPRUSRTM is a SAS time value ranging from 0 to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: GPRUSRTM contains the sum of all values of GPRUSRTM processed in the level of summarization.

USAGE NOTES: None

B.3 - The UNIX Application Performance (AUMAPP) File

The UNIX Application Performance file contains interval-based measurements useful for analyzing application use. It includes information about CPU use, disk I/Os, memory, transactions, processes, and wait states, and is derived from MeasureWare Application metrics.

The metrics supported vary by operating system.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	AUMAPU	YEAR	MONTH
	DAY	HOUR	ENDTS		
DAYS	SYSID	AUMOPSYS	AUMAPU	YEAR	MONTH
	DAY	HOUR			
WEEKS	SYSID	AUMOPSYS	AUMAPU	YEAR	WEEK
	ZONE	HOUR			
MONTHS	SYSID	AUMOPSYS	AUMAPU	YEAR	MONTH
	ZONE				
YEARS	SYSID	AUMOPSYS	AUMAPU	YEAR	ZONE
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
XDWMY.E	AUMAPU	- Application Name
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	SYSID	- System Identifier
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
		Common Data Elements
XDWMY.E	AUMAPNUM	- Application Number
XDWMY.E	AUMSAMPL	- Number of Data Samples
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
		Retained Data Elements
X.....E	APPPCSCP	- Pct CPU Wait
X.....E	APPPCSDK	- Pct Disk I/O Wait
X.....E	APPPCSDR	- Pct Time Directed to Wait
X.....E	APPPCSIM	- Pct Software Impede Wait
X.....E	APPPCSIP	- Pct Interprocess Communication Wait
X.....E	APPPCSLN	- Pct LAN I/O Wait
X.....E	APPPCSOT	- Pct Non-LAN/DISK/TERM I/O Wait
X.....E	APPPCSWP	- Pct Memory Wait
		Accumulated Data Elements
XDWMY.E	APPAPRCT	- Sum of Avg Processes Active
XDWMY.E	APPCOST	- Processing Charges
XDWMY.E	APPCPRCT	- Number of Processes Completed
XDWMY.E	APPCPUCM	- PRM CPU Cap Mode
XDWMY.E	APPCPUCP	- Total Available CPU Time
XDWMY.E	APPCPUEN	- PRM CPU Entitlement
XDWMY.E	APPCPUST	- PRM CPU State
XDWMY.E	APPCPUTM	- Total CPU Use
XDWMY.E	APPDBKCT	- Block I/Os
XDWMY.E	APPDBRCT	- Block Reads
XDWMY.E	APPDBWCT	- Block Writes
XDWMY.E	APPDFSCT	- Number of File System Reads and Writes
XDWMY.E	APPDMMCT	- Number of Memory Manager Reads and Write
XDWMY.E	APPDRAIO	- Number of Raw Reads and Writes
XDWMY.E	APPDSKIO	- Number of Physical Disk I/Os
XDWMY.E	APPDSYIO	- Number of System Reads and Writes
XDWMY.E	APPLIOCT	- Logical Disk I/Os
XDWMY.E	APPLOGMD	- PRM Logging Mode
XDWMY.E	APPLRDCT	- Logical Disk Reads
XDWMY.E	APPLWTCT	- Logical Disk Writes
XDWMY.E	APPMEMAV	- PRM Available Memory
XDWMY.E	APPMEMEN	- PRM MEM Entitlement
XDWMY.E	APPMEMPR	- Private Resident Memory Size in KBytes
XDWMY.E	APPMEMRS	- Resident Memory Size in KBytes
XDWMY.E	APPMEMST	- PRM MEM State
XDWMY.E	APPMEMSW	- Amount of Swap Space Configured
XDWMY.E	APPMEMSZ	- Amount of Main Memory Configured
XDWMY.E	APPMEMUP	- PRM MEM Upperbound
XDWMY.E	APPMEMUS	- PRM Memory Used by Processes

```
 XDWMY.E  APPMEMVR - Virtual Memory Size in KBytes
 XDWMY.E  APPNICTM - CPU Use at Nice Priorities
 XDWMY.E  APPNORTM - CPU Use at Normal Priorities
 XDWMY.E  APPPFMAJ - Page Faults to Disk
 XDWMY.E  APPPFMIN - Page Faults to Memory
 XDWMY.E  APPPKBCT - I/O Transfer Count in KBytes
 XDWMY.E  APPPRDCT - Physical Disk Reads
 XDWMY.E  APPPROCT - Sum of Avg Processes Alive
 XDWMY.E  APPPWTCT - Physical Disk Writes
 XDWMY.E  APPRTMTM - CPU Use at Real Time Priorities
 XDWMY.E  APPRUNTM - Total Run Time of Completing Processes
 XDWMY.E  APPSUSCT - Procs suspended during interval
 XDWMY.E  APPSYSTM - CPU Use for System Processing
 XDWMY.E  APPUSRTM - CPU Use in User Mode

Derived Data Elements

 XDWMY.E  APPAVAPR - Avg Processes Active
 XDWMY.E  APPAVPRO - Avg Processes Alive
 XDWMY.E  APPAVRTM - Avg Run Time of Completing Processes
 XDWMY.E  APPAVSUS - Avg Number of Processes Suspended
 XDWMY.E  APPDFSRT - Rate of File System Reads/Writes (IO/Sec
 XDWMY.E  APPDMMRT - Rate of Memory Mgr. Reads/Writes (IO/Sec
 XDWMY.E  APPDRART - Rate of Raw Reads/Writes (IO/Sec)
 XDWMY.E  APPDSYRT - Rate of System Reads/Writes (IO/Sec)
 XDWMY.E  APPMAJRT - Rate of Page Faults to Disk
 XDWMY.E  APPMINRT - Rate of Page Faults to Memory
 XDWMY.E  APPPCCPU - Percent Total CPU Use
 XDWMY.E  APPPCMPR - Percent Memory Use for Private Resident
 XDWMY.E  APPPCMRS - Percent Memory Use for Resident Memory
 XDWMY.E  APPPCNIC - Percent CPU Use at Nice Priorities
 XDWMY.E  APPPCNOR - Percent CPU Use at Normal Priorities
 XDWMY.E  APPPCRTM - Percent CPU Use at Real Time Priorities
 XDWMY.E  APPPCSYS - Percent CPU Use for System Processing
 XDWMY.E  APPPCUSR - Percent CPU Use in User Mode
 XDWMY.E  APPPKBRT - I/O Transfer Rate in KBytes
 XDWMY.E  APPPSPRD - Physical Disk Read Rate
 XDWMY.E  APPPSPWT - Physical Disk Write Rate


 +-----------------+
 | A P P A P R C T |  Sum of Avg Processes Active
 +-----------------+
 APPAPRCT contains the sum of the average active processes for
 this application during this measurement interval.

 TYPE:  Accumulated - Numeric, measurement unit is real
 number.

 SEE ALSO:  APPAVAPR - Avg Processes Active
            AUMSAMPL - Number of Data Samples
```

```
ELEMENT DERIVATION: At the DETAIL time-span, APPAPRCT is
computed using MeasureWare metrics in the following
calculation:

APP_ACTIVE_PROC*GBL_PROC_SAMPLE

where APP_ACTIVE_PROC is from Application metrics and
GBL_PROC_SAMPLE is from Global metrics.

RANGE OF VALUE:  APPAPRCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPAPRCT contains the sum of all
values of APPAPRCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P A V A P R |  Avg Processes Active
+-----------------+
APPAVAPR contains the average number of active processes
(consuming CPU time) for this application.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  APPAPRCT - Sum of Avg Processes Active
           AUMSAMPL - Number of Data Samples

ELEMENT DERIVATION:  APPAVAPR is computed using the following
algorithm:

APPAVAPR = AUMSAMPL;
IF AUMSAMPL
THEN APPAVAPR = APPAPRCT/AUMSAMPL;

RANGE OF VALUE:  APPAVAPR takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPAVAPR contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| A P P A V P R O |  Avg Processes Alive
+-----------------+
APPAVPRO contains the average number of processes alive in
the application.  A process that was alive for half of the
measurement interval counts as half a process.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  APPPROCT - Sum of Avg Processes Alive
           AUMSAMPL - Number of Data Samples

ELEMENT DERIVATION:  APPAVPRO is computed using the following
algorithm:

  APPAVPRO = AUMSAMPL;
  IF AUMSAMPL
  THEN APPAVPRO = APPPROCT/AUMSAMPL;

RANGE OF VALUE:  APPAVPRO takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPAVPRO contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P A V R T M |  Avg Run Time of Completing Processes
+-----------------+
APPAVRTM contains the average number of seconds of run time
accumulated by processes that completed during the
measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  APPRUNTM - Total Run Time of Completing Processes
           APPCPRCT - Total Processes Completed

ELEMENT DERIVATION:  APPAVRTM is computed using the following
algorithm:

  APPAVRTM = APPCPRCT;
  IF APPCPRCT
  THEN APPAVRTM = APPRUNTM/APPCPRCT;

RANGE OF VALUE:  APPAVRTM takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPAVRTM contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P A V S U S |  Avg Processes Suspended
+-----------------+
APPAVSUS contains the average number of suspended processes
for this application during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  APPSUSCT - Sum of Avg Processes Suspended
           AUMSAMPL - Number of Data Samples

ELEMENT DERIVATION:  APPAVSUS is computed using the following
algorithm:

  APPAVSUS = AUMSAMPL;
  IF AUMSAMPL
  THEN APPAVSUS = APPSUSCT/AUMSAMPL;

RANGE OF VALUE:  APPAVSUS takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPAVSUS contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P C O S T   |  Processing Charges
+-----------------+
APPCOST contains the processing charges accumulated for this
application during the measurement interval.  If this element
is selected for billing processing by CA MICS Accounting and
Chargeback, your user-written code will compute a value for
this element.

TYPE:  Accumulated - Numeric, measurement unit is real
number.
```

SEE ALSO: KPSCOST - Processing Charges

ELEMENT DERIVATION: APPCOST is set by code you write in the USRSAPP exit. If no value is set, this element's value is 0.

RANGE OF VALUE: APPCOST takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPCOST contains the sum of all values of APPCOST processed in the level of summarization.

USAGE NOTES: None

A P P C P R C T | Number of Processes Completed

APPCPRCT contains the number of processes belonging to this application that completed during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPRUNTM - Total Run Time of Completing Processes

ELEMENT DERIVATION: APPCPRCT is set to the value of the MeasureWare Application metric APP_COMPLETED_PROC.

RANGE OF VALUE: APPCPRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPCPRCT contains the sum of all values of APPCPRCT processed in the level of summarization.

USAGE NOTES: None

A P P C P U C M | PRM CPU Cap Mode

APPCPUCM contains a numeric variable that denotes one of the following meanings:

- 0 - PRM is not configured or installed
- 1 - CPU Cap Mode is not enabled (PRM CPU entitlements are in effect)
- 2 - CPU Cap mode is enabled (PRM CPU entitlements act as caps or limits)

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPLOGMD - PRM Logging Mode
APPMEMEN - PRM MEM Entitlement
APPMEMUP - PRM MEM Upperbound
APPMEMST - PRM MEM State
APPCPUST - PRM CPU State
APPMEMAV - PRM Available Memory
APPCPUEN - PRM CPU Entitlement
APPMEMUS - PRM Memory Used by Processes

ELEMENT DERIVATION: APPCPUCM is set to the value of the MeasureWare Application metric APP_PRM_CPUCAP_MODE.

RANGE OF VALUE: APPCPUCM ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPCPUCM contains the sum of all values of APPCPUCM processed in the level of summarization.

USAGE NOTES: None

A P P C P U C P | Total Available CPU Time

APPCPUCP contains the total amount of CPU time available to all applications during this measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: DURATION - Recording Interval Time
MCFNCPUS - Number of Processors Configured

ELEMENT DERIVATION: At the DETAIL time-span, APPCPUCP is computed using MeasureWare metrics in the following calculation:

INTERVAL*GBL_NUM_CPU

where INTERVAL is from Global data and GBL_NUM_CPU is from Configuration data.

RANGE OF VALUE: APPCPUCP is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPCPUCP contains the sum of all values of APPCPUCP processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| A P P C P U E N |  PRM CPU Entitlement
+-----------------+
APPCPUEN contains the Process Resource Manager (PRM) CPU
entitlement for a specific PRM Group Identifier entry.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  APPLOGMD - PRM Logging Mode
           APPMEMEN - PRM MEM Entitlement
           APPMEMUP - PRM MEM Upperbound
           APPMEMST - PRM MEM State
           APPCPUST - PRM CPU State
           APPMEMAV - PRM Available Memory
           APPCPUCM - PRM CPU Cap Mode
           APPMEMUS - PRM Memory Used by Processes

ELEMENT DERIVATION:  APPCPUEN is set to the value of the
MeasureWare Application metric APP_PRM_CPU_ENTITLEMENT.

RANGE OF VALUE:  APPCPUEN ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPCPUEN contains the sum of all
values of APPCPUEN processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P C P U S T |  PRM CPU State
+-----------------+
APPCPUST contains the state of the Performance Resource
Manager (PRM) CPU on a system.  The possible values for
this data elements are:

   0 - PRM is not installed
   1 - reset
   2 - configured/disabled
   3 - enabled

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  APPLOGMD - PRM Logging Mode
           APPMEMEN - PRM MEM Entitlement
           APPMEMUP - PRM MEM Upperbound
           APPMEMST - PRM MEM State
           APPMEMAV - PRM Available Memory
           APPCPUEN - PRM CPU Entitlement
           APPCPUCM - PRM CPU Cap Mode
           APPMEMUS - PRM Memory Used by Processes

ELEMENT DERIVATION:  APPCPUST is set to the value of
MeasureWare Application metric APP_PRM_STATE.

RANGE OF VALUE:  APPCPUST ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPCPUST contains the sum of all
values of APPCPUST processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P C P U T M |  Total CPU Use
+-----------------+
APPCPUTM contains the total amount of CPU time used by this
application during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  APPUSRTM - CPU Use in User Mode
           APPSYSTM - CPU Use for System Processing
           APPNICTM - CPU Use at Nice Priorities
           APPRTMTM - CPU Use at Real Time Priorities

ELEMENT DERIVATION:  APPCPUTM is set to the value of the
MeasureWare Application metric APP_CPU_TOTAL_TIME.

RANGE OF VALUE:  APPCPUTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPCPUTM contains the sum of all
values of APPCPUTM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P D B K C T |  Block I/Os
+-----------------+
APPDBKCT contains the number of block I/Os to the file system
cache for processes in this group during the measurement
interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  APPDBRCT - Block Reads
           APPDBWCT - Block Writes
```

ELEMENT DERIVATION: APPDBKCT is set to the value of the MeasureWare Application metric APP_DISK_BLOCK_IO.

RANGE OF VALUE: APPDBKCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDBKCT contains the sum of all values of APPDBKCT processed in the level of summarization.

USAGE NOTES: When a file is accessed, it is memory mapped by the operating system. Accesses generate virtual memory IOs. Reading a file generates block IOs as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM IOs) and posting updated inode information to disk (block IOs).

A P P D B R C T | Block Reads

APPDBRCT contains the number of block reads from the file system buffer cache for processes in this group during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPDBWCT - Block Writes
APPDBKCT - Block I/Os

ELEMENT DERIVATION: APPDBRCT is set to the value of the MeasureWare Application metric APP_DISK_BLOCK_READ.

RANGE OF VALUE: APPDBRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDBRCT contains the sum of all values of APPDBRCT processed in the level of summarization.

USAGE NOTES: When a file is accessed, it is memory mapped by the operating system. Accesses generate virtual memory IOs. Reading a file generates block IOs as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM IOs) and posting updated inode information to disk (block IOs).

A P P D B W C T | Block Writes

APPDBWCT contains the number of block writes to the file system buffer cache for processes in this group during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPDBRCT - Block Reads
APPDBKCT - Block I/Os

ELEMENT DERIVATION: APPDBWCT is set to the value of the MeasureWare Application metric APP_DISK_BLOCK_WRITE.

RANGE OF VALUE: APPDBWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDBWCT contains the sum of all values of APPDBWCT processed in the level of summarization.

USAGE NOTES: When a file is accessed, it is memory mapped by the operating system. Accesses generate virtual memory IOs. Reading a file generates block IOs as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM IOs) and posting updated inode information to disk (block IOs).

A P P D F S C T | Number of File System Reads and Writes

APPDFSCT contains the number of times the file system read from or wrote to a disk on behalf of processes assigned to this application during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPDMMCT - Number of Memory Mgr. Reads and Writes
APPDSYIO - Number of System Reads and Writes
APPDRAIO - Number of Raw Reads and Writes

ELEMENT DERIVATION: APPDFSCT is set to the value of the MeasureWare Application metric APP_DISK_FS_IO.

RANGE OF VALUE: APPDFSCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDFSCT contains the sum of all values of APPDFSCT processed in the level of summarization.

USAGE NOTES: None

APPDFSRT — Rate of File System Read/Writes (IO/Sec)

APPDFSRT contains the rate at which the File System read and wrote data on behalf of this application during this measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is integer.

SEE ALSO: APPDFSCT - Number of File System Reads and Writes
DURATION - Recording Interval Time

ELEMENT DERIVATION: APPDFSRT is computed using the following algorithm:

```
APPDFSRT = DURATION;
IF DURATION
THEN APPDFSRT = APPDFSCT/DURATION;
```

RANGE OF VALUE: APPDFSRT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDFSRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

APPDMMCT — Number of Memory Mgr. Reads and Writes

APPDMMCT contains the number of times the Memory Manager read from or wrote to a disk on behalf of processes assigned to this application during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPDFSCT - Number of File System Reads and Writes
APPDSYIO - Number of System Reads and Writes
APPDRAIO - Number of Raw Reads and Writes

ELEMENT DERIVATION: APPDMMCT is set to the value of the MeasureWare Application metric APP_DISK_VM_IO.

RANGE OF VALUE: APPDMMCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDMMCT contains the sum of all values of APPDMMCT processed in the level of summarization.

USAGE NOTES: None

APPDMMRT — Rate of Memory Mgr. Read/Writes (IO/Sec)

APPDMMRT contains the rate at which the Memory Manager read and wrote data on behalf of this application during this measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is integer.

SEE ALSO: APPDMMCT - Number of Memory Mgr. Reads and Writes
DURATION - Recording Interval Time

ELEMENT DERIVATION: APPDMMRT is computed using the following algorithm:

```
APPDMMRT = DURATION;
IF DURATION
THEN APPDMMRT = APPDMMCT/DURATION;
```

RANGE OF VALUE: APPDMMRT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDMMRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

APPDRAIO — Number of Raw Reads and Writes

APPDRAIO contains the number of raw reads and writes to disks on behalf of processes assigned to this application during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPDFSCT - Number of File System Reads and Writes
APPDMMCT - Number of Memory Mgr. Reads and Writes
APPDSYIO - Number of System Reads and Writes

ELEMENT DERIVATION: APPDRAIO is set to the value of the MeasureWare Application metric APP_DISK_RAW_IO.

RANGE OF VALUE: APPDRAIO takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDRAIO contains the sum of all values of APPDRAIO processed in the level of summarization.

USAGE NOTES: None

APPDRART — Rate of Raw Read/Writes (IO/Sec)

APPDRART contains the rate at which raw reads and writes occurred on behalf of this application during this measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is integer.

SEE ALSO: APPDRAIO - Number of Raw Reads and Writes
DURATION - Recording Interval Time

ELEMENT DERIVATION: APPDRART is computed using the following algorithm:

```
APPDRART = DURATION;
IF DURATION
THEN APPDRART = APPDRAIO/DURATION;
```

RANGE OF VALUE: APPDRART ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDRART contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

APPDSKIO — Number of Physical Disk I/Os

APPDSKIO contains the number of physical disk I/Os caused by processes assigned to this application during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPPRDCT - Physical Disk Reads
APPPWTCT - Physical Disk Writes

ELEMENT DERIVATION: APPDSKIO is set to the value of the MeasureWare Application metric APP_DISK_PHYS_IO.

RANGE OF VALUE: APPDSKIO takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDSKIO contains the sum of all values of APPDSKIO processed in the level of summarization.

USAGE NOTES: If APPDSKIO is unusually high, this application may not be efficiently accessing files.

APPDSYIO — Number of System Reads and Writes

APPDSYIO contains the number of times the system read from or wrote to a disk on behalf of processes assigned to this application during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPDFSCT - Number of File System Reads and Writes
APPDMMCT - Number of Memory Mgr. Reads and Writes
APPDRAIO - Number of Raw Reads and Writes

ELEMENT DERIVATION: APPDSYIO is set to the value of the MeasureWare Application metric APP_DISK_SYSTEM_IO.

RANGE OF VALUE: APPDSYIO takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDSYIO contains the sum of all values of APPDSYIO processed in the level of summarization.

USAGE NOTES: None

APPDSYRT — Rate of System Read/Writes (IO/Sec)

APPDSYRT contains the rate at which the System read and wrote data on behalf of this application during this measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is integer.

SEE ALSO: APPDSYIO - Number of System Reads and Writes
DURATION - Recording Interval Time

ELEMENT DERIVATION: APPDSYRT is computed using the following algorithm:

```
APPDSYRT = DURATION;
IF DURATION
THEN APPDSYRT = APPDSYIO/DURATION;
```

RANGE OF VALUE: APPDSYRT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPDSYRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

A P P L I O C T | Logical Disk I/Os

APPLIOCT contains the number of logical disk I/Os for this application during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPLRDCT - Logical Disk Reads
APPLWTCT - Logical Disk Writes

ELEMENT DERIVATION: APPLIOCT is set to the value of the MeasureWare Application metric APP_DISK_LOGL_IO.

RANGE OF VALUE: APPLIOCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPLIOCT contains the sum of all values of APPLIOCT processed in the level of summarization.

USAGE NOTES: None

A P P L O G M D | PRM Logging Mode

APPLOGMD contains the Performance Resource Manager (PRM) logging mode. There is a parm file setting read at the beginning of scopeux that can be used to override the default (which is zero (0)); whereby the logging mode is set to 1, signifying that PRM group data is being logged by the MeasureWare agent.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPMEMEN - PRM MEM Entitlement
APPMEMUP - PRM MEM Upperbound
APPMEMST - PRM MEM State
APPCPUST - PRM CPU State
APPMEMAV - PRM Available Memory
APPCPUEN - PRM CPU Entitlement
APPCPUCM - PRM CPU Cap Mode
APPMEMUS - PRM Memory Used by Processes

ELEMENT DERIVATION: APPLOGMD is set to the value of the MeasureWare Application metric APP_PRM_LOGGING_MODE.

RANGE OF VALUE: APPLOGMD ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPLOGMD contains the sum of all values of APPLOGMD processed in the level of summarization.

USAGE NOTES: None

A P P L R D C T | Logical Disk Reads

APPLRDCT contains the number of logical disk reads for this application during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPLWTCT - Logical Disk Writes

ELEMENT DERIVATION: APPLRDCT is set to the value of the MeasureWare Application metric APP_DISK_LOGL_READ.

RANGE OF VALUE: APPLRDCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPLRDCT contains the sum of all values of APPLRDCT processed in the level of summarization.

USAGE NOTES: None

APPLWTCT — Logical Disk Writes

APPLWTCT contains the number of logical disk writes for this application during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPLRDCT - Logical Disk Reads

ELEMENT DERIVATION: APPLWTCT is set to the value of the MeasureWare Application metric APP_DISK_LOGL_WRITE.

RANGE OF VALUE: APPLWTCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPLWTCT contains the sum of all values of APPLWTCT processed in the level of summarization.

USAGE NOTES: None

APPMAJRT — Rate of Page Faults to Disk

APPMAJRT contains the number of page faults per second that required a disk IO for processes in this group during the interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: APPPFMAJ - Page Faults to Disk

ELEMENT DERIVATION: APPMAJRT is computed using the following algorithm:

```
APPMAJRT = DURATION;
IF DURATION
THEN APPMAJRT = APPPFMAJ/DURATION;
```

RANGE OF VALUE: APPMAJRT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPMAJRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

APPMEMAV — PRM Available Memory

APPMEMAV is the amount of physical memory minus the amount of memory reserved for specific processes running in the PRM_SYS group 0. The value of APPMEMAV is dynamic and changes with system usage.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: APPMEMAV is set to the value of the MeasureWare Application metric APP_PRM_MEM_AVAIL.

RANGE OF VALUE: APPMEMAV ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPMEMAV contains the sum of all values of APPMEMAV processed in the level of summarization.

USAGE NOTES: None

APPMEMEN — PRM MEM Entitlement

APPMEMEN is defined in the Process Resource Manager (PRM) configuration file and is the PRM MEM entitlement for this PRM Group identifier.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO:
- APPLOGMD - PRM Logging Mode
- APPMEMUP - PRM MEM Upperbound
- APPMEMST - PRM MEM State
- APPCPUST - PRM CPU State
- APPMEMAV - PRM Available Memory
- APPCPUEN - PRM CPU Entitlement
- APPCPUCM - PRM CPU Cap Mode
- APPMEMUS - PRM Memory Used by Processes

ELEMENT DERIVATION: APPMEMEN is set to the value of the MeasureWare Application metric APP_PRM_MEM_ENTITLEMENT.

RANGE OF VALUE: APPMEMEN ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPMEMEN contains the sum of all values of APPMEMEN processed in the level of summarization.

USAGE NOTES: None

APPMEMPR — Private Resident Memory Size in KBytes

APPMEMPR contains the size of physical memory that was used as private resident memory for the application during the measurement interval. This includes process data and stack areas.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPMEMRS - Resident Memory Size in KBytes
APPMEMVR - Virtual Memory Size in KBytes

ELEMENT DERIVATION: At the DETAIL time-span, APPMEMPR is computed using MeasureWare metrics in the following calculation:

GBL_MEM_PHYS*APP_MEM_PRIVATE_UTIL

RANGE OF VALUE: APPMEMPR takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPMEMPR contains the sum of all values of APPMEMPR processed in the level of summarization.

USAGE NOTES: None

APPMEMRS — Resident Memory Size in KBytes

APPMEMRS contains the size of resident memory in KBytes for processes in this application during the measurement interval. This consists of text, data, stack, and the process' portion of shared memory.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPMEMPR - Private Resident Memory Size in KBytes
APPMEMVR - Virtual Memory Size in KBytes

ELEMENT DERIVATION: APPMEMRS is set to the value of the MeasureWare Application metric APP_MEM_RES.

RANGE OF VALUE: APPMEMRS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPMEMRS contains the sum of all values of APPMEMRS processed in the level of summarization.

USAGE NOTES: APPMEMRS may report higher than normal statistics for an application if a large number of processes are sharing memory regions.

APPMEMST — PRM MEM State

APPMEMST contains the Process Resource Manager (PRM) MEM state on the system. Possible values for APPMEMST are:

0 - PRM is not installed or there is no memory specification
1 - reset
2 - configured/disabled
3 - enabled

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPLOGMD - PRM Logging Mode
APPMEMEN - PRM MEM Entitlement
APPMEMUP - PRM MEM Upperbound
APPCPUST - PRM CPU State
APPMEMAV - PRM Available Memory
APPCPUEN - PRM CPU Entitlement
APPCPUCM - PRM CPU Cap Mode
APPMEMUS - PRM Memory Used by Processes

ELEMENT DERIVATION: APPMEMST is set to the value of the MeasureWare Application metric APP_PRM_MEM_STATE.

RANGE OF VALUE: APPMEMST ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPMEMST contains the sum of all values of APPMEMST processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| A P P M E M S W |  Amount of Swap Space Configured
+-----------------+
APPMEMSW contains the amount of swap space (in kilobytes)
configured for this system during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  MCFMEMSW - Swap Space Configured

ELEMENT DERIVATION:  APPMEMSW is set to the value of
MCFMEMSW.

RANGE OF VALUE:  APPMEMSW takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPMEMSW contains the sum of all
values of APPMEMSW processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P M E M S Z |  Amount of Main Memory Configured
+-----------------+
APPMEMSZ contains the amount of main memory (in kilobytes)
configured for this system during this measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  MCFMEMSZ - Main Memory Configured

ELEMENT DERIVATION:  APPMEMSZ is set to the value of
MCFMEMSZ.

RANGE OF VALUE:  APPMEMSZ takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPMEMSZ contains the sum of all
values of APPMEMSZ processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P M E M U P |  PRM MEM Upperbound
+-----------------+
APPMEMUP is defined in the Process Resource Manager (PRM)
configuration file and is the PRM MEM upperbound for this PRM
Group identifier.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  APPLOGMD - PRM Logging Mode
           APPMEMEN - PRM MEM Entitlement
           APPMEMUP - PRM MEM Upperbound
           APPMEMST - PRM MEM State
           APPCPUST - PRM CPU State
           APPMEMAV - PRM Available Memory
           APPCPUEN - PRM CPU Entitlement
           APPCPUCM - PRM CPU Cap Mode
           APPMEMUS - PRM Memory Used by Processes

ELEMENT DERIVATION:  APPMEMEN is set to the value of the
MeasureWare Application metric APP_PRM_MEM_UPPERBOUND.

RANGE OF VALUE:  APPMEMUP ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPMEMUP contains the sum of all
values of APPMEMUP processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P M E M U S |  PRM Memory Used by Processes
+-----------------+
APPMEMUS contains the amount of Process Resource Manager
(PRM) memory used by processes within the PRM groups
measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is real
number.

SEE ALSO:  APPMEMAV - PRM Available Memory

ELEMENT DERIVATION:  At the DETAIL timespan, APPMEMUS is
computed using the following calculation:

  (APP_PRM_MEM_UTIL*APPMEMAV)/100

where both APP_PRM_MEM_UTIL and APPMEMAV are from MeasureWare
Application data.

RANGE OF VALUE:  APPMEMUS takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPMEMUS contains the sum of all
values of APPMEMUS processed in the level of summarization.
```

USAGE NOTES: None

APPMEMVR — Virtual Memory Size in KBytes

APPMEMVR contains the size of virtual memory for processes in this application during the measurement interval. This is the sum of the virtual memory segment sizes for all processes in this application. Since virtual memory segment size includes shared memory, shared memory is counted multiple times.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPMEMPR - Private Resident Memory Size in KBytes
APPMEMRS - Resident Memory Size in KBytes

ELEMENT DERIVATION: APPMEMVR is set to the value of the MeasureWare Application metric APP_MEM_VIRT.

RANGE OF VALUE: APPMEMVR takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPMEMVR contains the sum of all values of APPMEMVR processed in the level of summarization.

USAGE NOTES: None

APPMINRT — Rate of Page Faults to Memory

APPMINRT contains the number of page faults per second satisfied in memory (that is, pages that were reclaimed from one of the free lists) for processes in this group during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: APPPFMIN - Page Faults to Memory

ELEMENT DERIVATION: APPMINRT is computed using the following algorithm:

```
APPMINRT = DURATION;
IF DURATION
THEN APPMINRT = APPPFMIN/DURATION;
```

RANGE OF VALUE: APPMINRT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPMINRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: Fluctuations in this average indicate that a user or application might make more efficient use of resources by maintaining its connection rather than logging onto and out of the system.

APPNICTM — CPU Use at Nice Priorities

APPNICTM contains the number of seconds that the CPU was used to run programs at a nice priority.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: APPPCNIC - Percent CPU Use at Nice Priorities
APPRTMTM - CPU Use at Real Time Priorities
APPSYSTM - CPU Use for System Processing
APPUSRTM - CPU Use in User Mode

ELEMENT DERIVATION: APPNICTM is set to the value of the MeasureWare Application metric APP_CPU_NICE_TIME.

RANGE OF VALUE: APPNICTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPNICTM contains the sum of all values of APPNICTM processed in the level of summarization.

USAGE NOTES: None

APPNORTM — CPU Use at Normal Priorities

APPNORTM contains the number of seconds that the CPU was used to run programs at normal priorities.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: APPNICTM - CPU Use at Nice Priorities
APPRTMTM - CPU Use at Real Time Priorities

ELEMENT DERIVATION: APPNORTM is set to the value of the MeasureWare Application metric APP_CPU_NORMAL_TIME.

RANGE OF VALUE: APPNORTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPNORTM contains the sum of all values of APPNORTM processed in the level of summarization.

USAGE NOTES: None

A P P P C C P U | Percent Total CPU Use

APPPCCPU contains the percentage of time that the CPU was used by this application during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: APPCPUCP - Total Available CPU Time
APPCPUTM - Total CPU Use

ELEMENT DERIVATION: APPPCCPU is computed using the following algorithm:

```
APPPCCPU = APPCPUCP;
IF APPCPUCP
THEN APPPCCPU = 100*APPCPUTM/APPCPUCP;
```

RANGE OF VALUE: APPPCCPU takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPCCPU contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

A P P P C M P R | Percent Memory Use for Private Resident

APPPCMPR contains the percentage of physical memory that was used as private resident memory for the application during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: APPMEMPR - Private Resident Memory Size in KBytes
APPMEMSZ - Amount of Main Memory Configured

ELEMENT DERIVATION: APPPCMPR is computed using the following algorithm:

```
APPPCMPR = APPMEMSZ;
IF APPMEMSZ
  THEN APPPCMPR = 100*APPMEMPR/APPMEMSZ;
```

RANGE OF VALUE: APPPCMPR takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: APPPCMPR contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

A P P P C M R S | Percent Memory Use for Resident Memory

APPPCMRS contains the percentage of physical memory that was used as resident memory for processes in this application during the measurement interval. This consists of text, data, stack, and the process' portion of shared memory.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: APPMEMRS - Resident Memory Size in KBytes
APPMEMSZ - Amount of Main Memory Configured

ELEMENT DERIVATION: APPPCMRS is computed using the following algorithm:

```
APPPCMRS = APPMEMSZ;
IF APPMEMSZ
THEN APPPCMRS = 100*APPMEMRS/APPMEMSZ;
```

RANGE OF VALUE: APPPCMRS takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: APPPCMRS contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

```
+-----------------+
| A P P P C N I C |  Percent CPU Use at Nice Priorities
+-----------------+
APPPCNIC contains the percentage of time that the CPU was
busy handling requests from this application whose priorities
had a NICE value.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  APPCPUCP - Total Available CPU Time
           APPNICTM - CPU Use at Nice Priorities

ELEMENT DERIVATION:  APPPCNIC is computed using the following
algorithm:

  APPPCNIC = APPCPUCP;
  IF APPCPUCP
  THEN APPPCNIC = 100*APPNICTM/APPCPUCP;

RANGE OF VALUE:  APPPCNIC takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPCNIC contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P P C N O R |  Percent CPU Use at Normal Priorities
+-----------------+
APPPCNOR contains the percentage of time that processes from
this application were in user mode running at normal priority
during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  APPCPUCP - Total Available CPU Time
           APPNORTM - CPU Use at Normal Priorities

ELEMENT DERIVATION:  APPPCNOR is computed using the following
algorithm:

  APPPCNOR = APPCPUCP;
  IF APPCPUCP
  THEN APPPCNOR = 100*APPNORTM/APPCPUCP;

RANGE OF VALUE:  APPPCNOR takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPCNOR contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P P C R T M |  Percent CPU Use at Real Time Priorities
+-----------------+
APPPCRTM contains the percentage of time that the CPU was
busy satisfying this application's requests whose priority
fell in the real time range.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  APPCPUCP - Total Available CPU Time
           APPRTMTM - CPU Use at Real Time Priorities

ELEMENT DERIVATION:  APPPCRTM is computed using the following
algorithm:

  APPPCRTM = APPCPUCP;
  IF APPCPUCP
  THEN APPPCRTM = 100*APPRTMTM/APPCPUCP;

RANGE OF VALUE:  APPPCRTM takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPCRTM contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P P C S C P |  Pct CPU Wait
+-----------------+
APPPCSCP contains the percentage of time processes belonging
to this application spend waiting for the CPU.

TYPE:  Retained - Numeric, measurement unit is real number.
```

SEE ALSO: Not applicable

ELEMENT DERIVATION: APPPCSCP is set to the value of the MeasureWare Application metric APP_PRI_WAIT_PCT.

RANGE OF VALUE: APPPCSCP takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPCSCP contains the last value of APPPCSCP processed in the level of summarization.

USAGE NOTES: If this application has more than 2 processes, APPPCSCP should be low. If the value is high, your system may have a CPU bottleneck.

A P P P C S D K | Pct Disk I/O Wait

APPPCSDK contains the percentage of time processes belonging to this application spent waiting for disk I/O.

TYPE: Retained - Numeric, measurement unit is real number.

SEE ALSO: APPPSPRD - Physical Disk Read Rate
APPPSPWT - Physical Disk Write Rate

ELEMENT DERIVATION: APPPCSDK is set to the value of the MeasureWare Application metric APP_DISK_SUBSYSTEM_WAIT_PCT.

RANGE OF VALUE: APPPCSDK takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPCSDK contains the last value of APPPCSDK processed in the level of summarization.

USAGE NOTES: If APPPCDSK's value is high, this application is I/O intensive. The system may benefit from tuning disk activity.

A P P P C S D R | Pct Time Directed to Wait

APPPCSDR contains the percentage of time that processes belonging to this application spent sleeping.

TYPE: Retained - Numeric, measurement unit is real number.

SEE ALSO: Not applicable

ELEMENT DERIVATION: APPPCSDR is set to the value of the MeasureWare Application metric APP_SLEEP_WAIT_PCT.

RANGE OF VALUE: APPPCSDR takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPCSDR contains the last value of APPPCSDR processed in the level of summarization.

USAGE NOTES: The sleep state does not indicate a resource bottleneck.

A P P P C S I M | Pct Software Impede Wait

APPPCSIM contains the percentage of time processes belonging to this application spent waiting due to software impedes.

TYPE: Retained - Numeric, measurement unit is real number.

SEE ALSO: Not applicable

ELEMENT DERIVATION: APPPCSIM is set to the value of the MeasureWare Application metric APP_SEM_WAIT_PCT.

RANGE OF VALUE: APPPCSIM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPCSIM contains the last value of APPPCSIM processed in the level of summarization.

USAGE NOTES: None

A P P P C S I P | Pct Interprocess Communication Wait

APPPCSIP contains the percentage of time processes belonging to this application spend waiting for interprocess communications.

TYPE: Retained - Numeric, measurement unit is real number.

SEE ALSO: Not applicable

ELEMENT DERIVATION: APPPCSIP is set to the value of the MeasureWare Application metric APP_IPC_SUBSYSTEM_WAIT_PCT.

RANGE OF VALUE: The value of APPPCSIP ranges from a minimum of 0 to a maximum of 100.

SUMMARIZATION PROCESS: APPPCSIP contains the last value of APPPCSIP processed in the level of summarization.

USAGE NOTES: None

| A P P P C S L N | Pct LAN I/O Wait

APPPCSLN contains the percentage of time that processes belonging to this application spent waiting for LAN I/O.

TYPE: Retained - Numeric, measurement unit is real number.

SEE ALSO: Not applicable

ELEMENT DERIVATION: APPPCSLN is set to the value of the MeasureWare Application metric APP_NETWORK_SUBSYSTEM_WAIT_PCT.

RANGE OF VALUE: APPPCSLN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPCSLN contains the last value of APPPCSLN processed in the level of summarization.

USAGE NOTES: None

| A P P P C S O T | Pct Non-LAN/DISK/TERM I/O Wait

APPPCSOT contains the percentage of time that processes belonging to this application spent waiting for I/O from devices other than LANs, disks, and terminals.

TYPE: Retained - Numeric, measurement unit is real number.

SEE ALSO: APPPCSDK - Percent Disk I/O Wait
APPPCSLN - Percent LAN I/O Time

ELEMENT DERIVATION: APPPCSOT is set to the value of the MeasureWare Application metric APP_OTHER_IO_WAIT_PCT.

RANGE OF VALUE: APPPCSOT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPCSOT contains the last value of APPPCSOT processed in the level of summarization.

USAGE NOTES: None

| A P P P C S W P | Pct Memory Wait

APPPCSWP contains the percentage of time processes belonging to this application are waiting for memory.

TYPE: Retained - Numeric, measurement unit is real number.

SEE ALSO: APPMEMSZ - Amount of Main Memory Configured

ELEMENT DERIVATION: APPPCSWP is set to the value of the MeasureWare Application metric APP_MEM_WAIT_PCT.

RANGE OF VALUE: APPPCSWP takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPCSWP contains the last value of APPPCSWP processed in the level of summarization.

USAGE NOTES: If APPPCSWP's value is high, your system may have a memory bottleneck or be suffering buffer cache contention.

| A P P P C S Y S | Percent CPU Use for System Processing

APPPCSYS contains the percentage of CPU used by this application in system mode, for context switching, and for interrupt handling during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: APPCPUCP - Total Available CPU Time
APPSYSTM - CPU Use for System Processing

ELEMENT DERIVATION: APPPCSYS is computed using the following algorithm:

```
APPPCSYS = APPCPUCP;
IF APPCPUCP
THEN APPPCSYS = 100*APPSYSTM/APPCPUCP;
```

RANGE OF VALUE: APPPCSYS takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: APPPCSYS contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

A P P P C U S R | Percent CPU Use in User Mode

APPPCUSR contains the amount of time that this application spent executing in user mode.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: APPCPUCP - Total Available CPU Time
APPUSRTM - CPU Use in User Mode

ELEMENT DERIVATION: APPPCUSR is computed using the following algorithm:

```
APPPCUSR = APPCPUCP;
IF APPCPUCP
THEN APPPCUSR = 100*APPUSRTM/APPCPUCP;
```

RANGE OF VALUE: APPPCUSR takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: APPPCUSR contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

A P P P F M A J | Page Faults to Disk

APPPFMAJ contains the number of page faults that required a disk IO for processes in this group during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: APPMAJRT - Rate of Page Faults to Disk

ELEMENT DERIVATION: At the DETAIL time-span, APPPFMAJ is computed using MeasureWare Global metrics in the following calculation:

APP_MAJOR_FAULT_RATE*INTERVAL

RANGE OF VALUE: APPPFMAJ takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPFMAJ contains the sum of all values of APPPFMAJ processed in the level of summarization.

USAGE NOTES: None

A P P P F M I N | Page Faults to Memory

APPPFMIN contains the number of page faults satisfied in memory (that is, pages that were reclaimed from one of the free lists) for processes in this group during the interval.

TYPE: Accumulated - Numeric, measurement unit is real integer.

SEE ALSO: APPMINRT - Rate of Page Faults to Memory

ELEMENT DERIVATION: At the DETAIL time-span, APPPFMIN is computed using MeasureWare Global metrics in the following calculation:

APP_MINOR_FAULT_RATE*INTERVAL

RANGE OF VALUE: APPPFMIN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPPFMIN contains the sum of all values of APPPFMIN processed in the level of summarization.

USAGE NOTES: None

A P P P K B C T | I/O Transfer Count in KBytes

APPPKBCT contains the number of characters transferred to all devices for this application during the measurement interval. The number is expressed in kilobytes.

TYPE: Accumulated - Numeric, measurement unit is integer.

```
SEE ALSO:  APPDSKIO - Number of Physical Disk I/Os
           APPPRDCT - Physical Disk Reads

ELEMENT DERIVATION:  APPPKBCT is set to the value of the
MeasureWare Application metric APP_IO_BYTE.

RANGE OF VALUE:  APPPKBCT ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPKBCT contains the sum of all
values of APPPKBCT processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| A P P P K B R T |  I/O Transfer Rate in KBytes
+-----------------+
APPPKBRT contains the I/O transfer rate, in kilobytes, for
this application during this measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is integer.

SEE ALSO:  APPPKBCT - I/O Transfer Count in KBytes
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  APPPKBRT is computed using the following
algorithm:

  APPPKBRT = DURATION;
  IF DURATION
  THEN APPPKBRT = APPPKBCT/DURATION;

RANGE OF VALUE:  APPPKBRT ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPKBRT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| A P P P R D C T |  Physical Disk Reads
+-----------------+
APPPRDCT contains the number of physical disk reads for this
application during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  APPDSKIO - Number of Physical Disk I/Os
           APPPWTCT - Physical Disk Writes

ELEMENT DERIVATION:  APPPRDCT is set to the value of the
MeasureWare Application metric APP_DISK_PHYS_READ.

RANGE OF VALUE:  APPPRDCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPRDCT contains the sum of all
values of APPPRDCT processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| A P P P R O C T |  Sum of Avg Processes Alive
+-----------------+
APPPROCT contains the average number of processes for this
application that are alive during this measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is real
number.

SEE ALSO:  APPAVPRO - Avg Processes Alive
           AUMSAMPL - Number of Samples

ELEMENT DERIVATION:  At the DETAIL time-span, APPPROCT is
computed using MeasureWare metrics in the following
calculation:

APP_ALIVE_PROC*GBL_PROC_SAMPLE

where APP_ALIVE_PROC is from Application metrics and
GBL_PROC_SAMPLE is from Global metrics.

RANGE OF VALUE:  APPPROCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPROCT contains the sum of all
values of APPPROCT processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| A P P P S P R D |  Physical Disk Read Rate
+-----------------+
APPPSPRD contains the rate at which physical disk reads
occurred for this application.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           APPPRDCT - Physical Disk Reads

ELEMENT DERIVATION:  APPPSPRD is computed using the following
algorithm:

  APPPSPRD = DURATION;
  IF DURATION
  THEN APPPSPRD = APPPRDCT/DURATION;

RANGE OF VALUE:  APPPSPRD takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPSPRD contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P P S P W T |  Physical Disk Write Rate
+-----------------+
APPPSPWT contains the rate at which physical disk writes
occurred for this application.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  APPPWTCT - Physical Disk Writes
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  APPPSPWT is computed using the following
algorithm:

  APPPSPWT = DURATION;
  IF DURATION
  THEN APPPSPWT = APPPWTCT/DURATION;

RANGE OF VALUE:  APPPSPWT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPSPWT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P P W T C T |  Physical Disk Writes
+-----------------+
APPPWTCT contains the number of physical disk writes for this
application during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  APPDSKIO - Number of Physical Disk I/Os
           APPPRDCT - Physical Disk Reads

ELEMENT DERIVATION:  APPPWTCT is set to the value of the
MeasureWare Application metric APP_DISK_PHYS_WRITE.

RANGE OF VALUE:  APPPWTCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPPWTCT contains the sum of all
values of APPPWTCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| A P P R T M T M |  CPU Use at Real Time Priorities
+-----------------+
APPRTMTM contains the number of seconds that the CPU was
used to run programs at real time priorities.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  APPNICTM - CPU Use at Nice Priorities
           APPPCRTM - Percent CPU Use at Real Time Priorities
           APPSYSTM - CPU Use for System Processing
           APPUSRTM - CPU Use in User Mode

ELEMENT DERIVATION:  APPRTMTM is set to the value of the
MeasureWare Application metric APP_CPU_REALTIME_TIME.
```

RANGE OF VALUE: APPRTMTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPRTMTM contains the sum of all values of APPRTMTM processed in the level of summarization.

USAGE NOTES: None

APPRUNTM — Total Run Time of Completing Processes

APPRUNTM contains the total run time for processes that completed for this application during this measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: APPCPRCT - Number of Processes Completed

ELEMENT DERIVATION: At the DETAIL time-span, APPRUNTM is computed using MeasureWare Application metrics in the following calculation:

APP_PROC_RUN_TIME*APP_COMPLETED_PROC

RANGE OF VALUE: APPRUNTM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPRUNTM contains the sum of all values of APPRUNTM processed in the level of summarization.

USAGE NOTES: None

APPSUSCT — Sum of Avg Processes Active

APPSUSCT contains the sum of the average suspended processes for this application during this measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: APPAVSUS - Avg Processes Suspended
AUMSAMPL - Number of Data Samples

ELEMENT DERIVATION: At the DETAIL time-span, APPSUSCT is computed using MeasureWare metrics in the following calculation:

APP_SUSPENDED_PROCS*GBL_PROC_SAMPLE

where APP_SUSPENDED_PROC is from Application metrics and GBL_PROC_SAMPLE is from Global metrics.

RANGE OF VALUE: APPSUSCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPSUSCT contains the sum of all values of APPSUSCT processed in the level of summarization.

USAGE NOTES: None

APPSYSTM — CPU Use for System Processing

APPSYSTM contains the number of seconds that the CPU was used for system processing, including time spent in system mode, for context switching, and for interrupt handling, during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: APPNICTM - CPU Use at Nice Priorities
APPPCSYS - Percent CPU Use for System Processing
APPRTMTM - CPU Use at Real Time Priorities
APPUSRTM - CPU Use in User Mode

ELEMENT DERIVATION: APPSYSTM is set to the value of the MeasureWare Application metric APP_CPU_SYS_MODE_TIME.

RANGE OF VALUE: APPSYSTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: APPSYSTM contains the sum of all values of APPSYSTM processed in the level of summarization.

USAGE NOTES: None

APPUSRTM — CPU Use in User Mode

APPUSRTM contains the number of seconds that the CPU was used in user mode during the measurement interval.

```
TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  APPNICTM - CPU Use at Nice Priorities
           APPPCUSR - Percent CPU Use in User Mode
           APPRTMTM - CPU Use at Real Time Priorities
           APPSYSTM - CPU Use for System Processing

ELEMENT DERIVATION:  At the DETAIL time-span, APPUSRTM is
computed using MeasureWare metrics for HP-UX systems in the
following calculation:

APP_CPU_NORMAL_TIME+APP_CPU_NICE_TIME+APP_CPU_REALTIME_TIME

RANGE OF VALUE:  APPUSRTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  APPUSRTM contains the sum of all
values of APPUSRTM processed in the level of summarization.

USAGE NOTES:  None
```

B.4 - The UNIX Process Activity (AUMKPS) File

The UNIX Process Activity file contains measurements about completed processes and is useful for accounting and process analysis. It includes information about CPU, disk, memory, and transactions, and is derived from MeasureWare Application metrics.

The metrics supported vary by operating system.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID YEAR	AUMOPSYS MONTH	AUMACT1 DAY	AUMACT2 HOUR	AUMACT3 ENDTS
DAYS	N/A				
WEEKS	N/A				
MONTHS	SYSID MONTH	AUMOPSYS ZONE	AUMACT1	AUMACT2	YEAR
YEARS	SYSID ZONE	AUMOPSYS	AUMACT1	AUMACT2	YEAR
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
X..MY.E	AUMACT1	- DIVISION
X..MY.E	AUMACT2	- USER GROUP
X.....E	AUMACT3	- USER NAME
X..MY.E	AUMOPSYS	- Operating System Name
X.....E	DAY	- Day of Month
X.....E	HOUR	- Hour of Day
X..M..E	MONTH	- Month of Year
X..MY.E	SYSID	- System Identifier
X.....E	WEEK	- Week of Year
X..MY.E	YEAR	- Year of Century
X..MY.E	ZONE	- Time Zone
		Common Data Elements
X.....E	AUMAPNUM	- Application Number
X.....E	AUMAPU	- Application Name
X.....E	AUMGRPID	- Group Identification Number
X.....E	AUMLNTTY	- Logon TTY Name
X.....E	AUMPIN	- Process Identification Number
X.....E	AUMPPID	- Parent Identification Number
X.....E	AUMPRI	- Dispatch Priority
X.....E	AUMPROG	- Program Name
X.....E	AUMUSER	- User Logon Name
X.....E	DAYNAME	- Name of Day of Week
X..MY.E	DURATION	- Recording Interval Time
X..MY.E	ENDTS	- End Time Stamp
X..MY.E	STARTTS	- Start Time Stamp
		Accumulated Data Elements
X..MY.E	KPSCOST	- Processing Charges
X..MY.E	KPSCPUTM	- Total CPU Use
X..MY.E	KPSDBKCT	- Block I/Os
X..MY.E	KPSDSKIO	- Physical Disk I/Os
X..MY.E	KPSLDKCT	- Logical Disk I/Os
X..MY.E	KPSPKBCT	- Total I/O Transfer Count in KBytes
X..MY.E	KPSRSSKB	- Resident Set Size
X..MY.E	KPSRUNTM	- Total Run Time
X..MY.E	KPSSCPTM	- CPU Wait Time
X..MY.E	KPSSDKTM	- Disk I/O Wait Time
X..MY.E	KPSSDRTM	- Directed to Wait Time
X..MY.E	KPSSIMTM	- Software Impede Wait Time
X..MY.E	KPSSIPTM	- IPC Resource Wait Time
X..MY.E	KPSSLNTM	- LAN I/O Wait Time
X..MY.E	KPSSNFTM	- Network File System Wait Time
X..MY.E	KPSSOITM	- Non-LAN/DISK/TERM I/O Wait Time
X..MY.E	KPSSSYTM	- System Wait Time
X..MY.E	KPSSVMTM	- Virtual Memory Wait Time
X..MY.E	KPSVSSKB	- Virtual Memory Regions
		Derived Data Elements
X..MY.E	KPSAVDRT	- Physical Disk I/O Rate
X..MY.E	KPSAVLRT	- Logical Disk I/O Rate
X..MY.E	KPSDBKRT	- Block I/O Rate
X..MY.E	KPSPKBRT	- I/O Transfer Rate in KBytes

```
+-----------------+
| K P S A V D R T |  Physical Disk I/O Rate
+-----------------+
KPSAVDRT contains the physical disk I/O rate for this
completed process.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           KPSDSKIO - Physical Disk I/Os

ELEMENT DERIVATION:  KPSAVDRT is computed using the following
algorithm:

  KPSAVDRT = DURATION;
  IF DURATION
    THEN KPSAVDRT = KPSDSKIO/DURATION;

RANGE OF VALUE:  KPSAVDRT takes any numeric value ranging
from 0 to a maximum that is limited only by practical
bounds.

SUMMARIZATION PROCESS:  KPSAVDRT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| K P S A V L R T |  Logical Disk I/O Rate
+-----------------+
KPSAVLRT contains the logical disk I/O rate for this
completed process.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           KPSLDKCT - Logical Disk I/Os

ELEMENT DERIVATION:  KPSAVLRT is computed using the following
algorithm:

  KPSAVLRT = DURATION;
  IF DURATION
    THEN KPSAVLRT = KPSLDKCT/DURATION;

RANGE OF VALUE:  KPSAVLRT takes any numeric value ranging
from 0 to a maximum that is limited only by practical
bounds.

SUMMARIZATION PROCESS:  KPSAVLRT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| K P S C O S T   |  Processing Charges
+-----------------+
KPSCOST contains the processing charges accumulated by this
completed process.  If this element is selected for billing
processing by CA MICS Accounting and Chargeback, your
user-written code will compute a value for this element.

TYPE:  Accumulated - Numeric, measurement unit is real
number.

SEE ALSO:  APPCOST - Processing Charges

ELEMENT DERIVATION:  KPSCOST is set by code you write in the
USRSKPS exit.  If no value is set, this element's value is 0.

RANGE OF VALUE:  KPSCOST takes any numeric value ranging from
0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  KPSCOST contains the sum of all
values of KPSCOST processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| K P S C P U T M |  Total CPU Use
+-----------------+
KPSCPUTM contains the total amount of CPU time used by
this completed process.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  KPSRUNTM - Total Run Time

ELEMENT DERIVATION:  KPSCPUTM is set to the value of the
MeasureWare Process metric PROC_CPU_TOTAL_TIME_CUM.

RANGE OF VALUE:  KPSCPUTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.
```

SUMMARIZATION PROCESS: KPSCPUTM contains the sum of all values of KPSCPUTM processed in the level of summarization.

USAGE NOTES: KPSCPUTM has a precision of one-hundredth (0.01) of a second.

K P S D B K C T | Block I/Os

KPSDBKCT contains the number of block I/Os made by or for a process during its lifetime or over the cumulative measurement interval. These are physical IOs generated by file system access and do not include virtual memory IO, or IOs relating to raw disk access. These are IOs for inode and superblock updates that are handled through the buffer cache.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDBRCT - Block Reads
IPSDBWCT - Block Writes

ELEMENT DERIVATION: KPSDBKCT is set to the value of the MeasureWare Application metric PROC_DISK_BLOCK_IO_CUM.

RANGE OF VALUE: KPSDBKCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSDBKCT contains the sum of all values of KPSDBKCT processed in the level of summarization.

USAGE NOTES: When a file is accessed, it is memory mapped by the operating system. Accesses generate virtual memory IOs. Reading a file generates block IOs as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM IOs) and posting updated inode information to disk (block IOs).

K P S D B K R T | Block I/O Rate

KPSDBRCT contains the logical read rate in IOs per second during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: KPSDBKCT - Block I/Os

DURATION - Recording Interval Time

ELEMENT DERIVATION: DDARKBRT is computed using the following algorithm:

```
KPSDBKRT = DURATION;
IF DURATION
  THEN KPSDBKRT = KPSDBKCT/DURATION;
```

RANGE OF VALUE: DDARKBRT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDARKBRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

K P S D S K I O | Physical Disk I/Os

KPSDSKIO contains the number of physical disk I/Os accumulated by this completed process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: KPSLDKCT - Logical Disk I/Os

ELEMENT DERIVATION: KPSDSKIO is set to the value of the MeasureWare Process metric PROC_DISK_PHYS_IO_CUM.

RANGE OF VALUE: KPSDSKIO takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSDSKIO contains the sum of all values of KPSDSKIO processed in the level of summarization.

USAGE NOTES: None

K P S L D K C T | Logical Disk I/Os

KPSLDKCT contains the number of logical disk I/Os requested by this completed process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: KPSDSKIO - Physical Disk I/Os

ELEMENT DERIVATION: KPSLDKCT is set to the value of the MeasureWare Process metric PROC_DISK_LOGL_IO_CUM.

RANGE OF VALUE: KPSLDKCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSLDKCT contains the sum of all values of KPSLDKCT processed in the level of summarization.

USAGE NOTES: None

K P S P K B C T | Total I/O Transfer Count in KBytes

KPSPKBCT contains the number of kilobytes of data read or written by a process during its lifetime.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: KPSDSKIO - Physical Disk I/Os

ELEMENT DERIVATION: KPSPKBCT is set to the value of the MeasureWare Process metric PROC_IO_BYTE_CUM.

RANGE OF VALUE: KPSPKBCT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSPKBCT contains the sum of all values of KPSPKBCT processed in the level of summarization.

USAGE NOTES: None

K P S P K B R T | I/O Transfer Rate in KBytes

KPSPKBRT contains the I/O transfer rate, in kilobytes, for this process during its lifetime.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is integer.

SEE ALSO: DURATION - Recording Interval Time
KPSPKBCT - Total I/O Transfer Count in KBytes

ELEMENT DERIVATION: KPSPKBRT is computed using the following algorithm:

```
KPSPKBRT = DURATION;
IF DURATION
THEN KPSPKBRT = KPSPKBCT/DURATION;
```

RANGE OF VALUE: KPSPKBRT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSPKBRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

K P S R S S K B | Resident Set Size

KPSRSSKB contains the resident set size, in kilobytes, for this completed process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSRSSKB - Resident Set Size

ELEMENT DERIVATION: KPSRSSKB is set to the value of IPSRSSKB.

RANGE OF VALUE: KPSRSSKB takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSRSSKB contains the sum of all values of KPSRSSKB processed in the level of summarization.

USAGE NOTES: None

K P S R U N T M | Total Run Time

KPSRUNTM contains the total amount of run time accumulated by this completed process.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSRUNTM - Total Accumulated Run Time

ELEMENT DERIVATION: KPSRUNTM is set to the value of

IPSRUNTM.

RANGE OF VALUE: KPSRUNTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSRUNTM contains the sum of all values of KPSRUNTM processed in the level of summarization.

USAGE NOTES: None

| K P S S C P T M | CPU Wait Time

KPSSCPTM contains the total amount of time the process waited for the CPU during its lifetime.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: KPSSCPTM is set to the value of the MeasureWare Process metric PROC_PRI_WAIT_TIME.

RANGE OF VALUE: KPSSCPTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSSCPTM contains the sum of all values of KPSSCPTM processed in the level of summarization.

USAGE NOTES: None

| K P S S D K T M | Disk I/O Wait Time

KPSSDKTM contains the disk I/O wait time accumulated by this completed process.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSSDKTM - Time Waited for Disk

ELEMENT DERIVATION: KPSSDKTM is set to the value of IPSSDKTM.

RANGE OF VALUE: KPSSDKTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSSDKTM contains the sum of all values of KPSSDKTM processed in the level of summarization.

USAGE NOTES: None

| K P S S D R T M | Directed to Wait Time

KPSSDRTM contains the total amount of time the process was directed to wait during its lifetime.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: KPSSDRTM is set to the value of the MeasureWare Process metric PROC_SLEEP_WAIT_TIME.

RANGE OF VALUE: KPSSDRTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSSDRTM contains the sum of all values of KPSSDRTM processed in the level of summarization.

USAGE NOTES: None

| K P S S I M T M | Software Impede Wait Time

KPSSIMTM contains the number of seconds that the process waited for software impedes during its lifetime.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: KPSSIMTM is set to the value of set MeasureWare Process metric PROC_SEM_WAIT_TIME.

RANGE OF VALUE: KPSSIMTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSSIMTM contains the sum of all values of KPSSIMTM processed in the level of summarization.

USAGE NOTES: None

| K P S S I P T M | IPC Resource Wait Time

KPSSIPTM contains the number of seconds that the process waited for InterProcess Communications during its lifetime.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: KPSSIPTM is set to the value of the MeasureWare Process metric PROC_IPC_SUBSYSTEM_WAIT_TIME.

RANGE OF VALUE: KPSSIPTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSSIPTM contains the sum of all values of KPSSIPTM processed in the level of summarization.

USAGE NOTES: None

| K P S S L N T M | LAN I/O Wait Time

KPSSLNTM contains the number of seconds that the process waited for LAN I/O during its lifetime.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: KPSSLNTM is set to the value of the MeasureWare Process metric PROC_LAN_WAIT_TIME.

RANGE OF VALUE: KPSSLNTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSSLNTM contains the sum of all values of KPSSLNTM processed in the level of summarization.

USAGE NOTES: None

| K P S S N F T M | Network File System Wait Time

KPSSNFTM contains the number of seconds that the process waited for the Network File System during its lifetime.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: KPSSNFTM is set to the value of the MeasureWare Process metric PROC_NFS_WAIT_TIME.

RANGE OF VALUE: KPSSNFTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSSNFTM contains the sum of all values of KPSSNFTM processed in the level of summarization.

USAGE NOTES: None

| K P S S O I T M | Non-LAN/DISK/TERM I/O Wait Time

KPSSOITM contains the number of seconds that the process waited for I/O that is not specifically tracked by individual elements during its lifetime.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: Not applicable

ELEMENT DERIVATION: KPSSOITM is set to the value of the MeasureWare Process metric PROC_OTHER_IO_WAIT_TIME.

RANGE OF VALUE: KPSSOITM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: KPSSOITM contains the sum of all values of KPSSOITM processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| K P S S S Y T M |  System Wait Time
+-----------------+
KPSSSYTM contains the number of seconds that the process
waited for the system during its lifetime.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  KPSSSYTM is set to the value of the
MeasureWare Process metric PROC_SYS_WAIT_TIME.

RANGE OF VALUE:  KPSSSYTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  KPSSSYTM contains the sum of all
values of KPSSSYTM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| K P S S V M T M |  Virtual Memory Wait Time
+-----------------+
KPSSVMTM contains the number of seconds that the process
waited for virtual memory I/O during its lifetime.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  KPSSVMTM is set to the value of the
MeasureWare Process metric PROC_MEM_WAIT_TIME.

RANGE OF VALUE:  KPSSVMTM is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  KPSSVMTM contains the sum of all
values of KPSSVMTM processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| K P S V S S K B |  Virtual Memory Regions
+-----------------+
KPSVSSKB contains the number of kilobytes of virtual memory
reserved for this completed process.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  IPSVSSKB - Virtual Memory Regions

ELEMENT DERIVATION:  KPSVSSKB is set to the value of
IPSVSSKB.

RANGE OF VALUE:  KPSVSSKB takes any numeric value ranging
from 0 to a maximum that is limited only by practical
bounds.

SUMMARIZATION PROCESS:  KPSVSSKB contains the sum of all
values of KPSVSSKB processed in the level of summarization.

USAGE NOTES:  None
```

B.5 - The UNIX Process Exception (AUMIPS) File

The UNIX Process Exception file contains measurements about interesting processes (processes that exceed thresholds) and is useful for performance and problem analyses and modeling. It includes information about CPU use, disk, memory, and transactions as well as reason codes (defining why the process is interesting). The AUMIPS file is derived from MeasureWare Application metrics.

The metrics supported vary by operating system.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity
DETAIL	SYSID AUMOPSYS YEAR MONTH DAY HOUR ENDTS AUMUSER AUMLNTTY AUMPIN AUMPROG
DAYS	N/A
WEEKS	N/A
MONTHS	N/A
YEARS	N/A
TABLES	N/A

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

```
 Time-    Data      Data Element
 Span *  Element    Description (LABEL)
 -------  --------   ----------------------------------------
```

Sequence/Summary Data Elements

```
X.....E  AUMOPSYS - Operating System Name
X.....E  AUMPIN   - Process Identification Number
X.....E  DAY      - Day of Month
X.....E  HOUR     - Hour of Day
X.....E  SYSID    - System Identifier
X.....E  WEEK     - Week of Year
X.....E  YEAR     - Year of Century
X.....E  ZONE     - Time Zone
```

Common Data Elements

```
X.....E  AUMAPNUM - Application Number
X.....E  AUMAPU   - Application Name
X.....E  AUMGRPID - Group Identification Number
X.....E  AUMLNTTY - Logon TTY Name
X.....E  AUMPPID  - Parent Identification Number
X.....E  AUMPRI   - Dispatch Priority
X.....E  AUMPROG  - Program Name
X.....E  AUMUSER  - User Logon Name
X.....E  DAYNAME  - Name of Day of Week
X.....E  DURATION - Recording Interval Time
X.....E  ENDTS    - End Time Stamp
X.....E  MONTH    - Month of Year
X.....E  STARTTS  - Start Time Stamp
```

Retained Data Elements

```
X.....E  IPSINTR  - Interest Reason Codes
X.....E  IPSPRMID - PRM Group Identification Number
X.....E  IPSSTOPR - Blocked State
```

Accumulated Data Elements

```
X.....E  IPSALITM - Process Alive Time
X.....E  IPSCPUTM - Total CPU Use
X.....E  IPSCSWTM - CPU Use for Context Switching
X.....E  IPSDBKCT - Block I/Os
X.....E  IPSDBRCT - Block Reads
X.....E  IPSDBWCT - Block Writes
X.....E  IPSDFRCT - File System Disk Reads
X.....E  IPSDFWCT - File System Disk Writes
X.....E  IPSDMMCT - Virtual Memory I/Os
X.....E  IPSDSKIO - Physical Disk I/Os
X.....E  IPSDSYCT - File System Management I/Os
X.....E  IPSFSDCT - File System Disk I/Os
X.....E  IPSINTTM - CPU Use for Interrupt Handling
X.....E  IPSLRDCT - Logical Disk Reads
X.....E  IPSLWTCT - Logical Disk Writes
X.....E  IPSNICTM - CPU Use at Nice Priority
X.....E  IPSNORTM - CPU Use at Normal Priority
X.....E  IPSPFMAJ - Page Faults to Disk
X.....E  IPSPFMIN - Page Faults to Memory
X.....E  IPSPKBCT - I/O Transfer Count in KBytes
X.....E  IPSRSSKB - Resident Set Size
X.....E  IPSRTMTM - CPU Use at Real Time Priority
X.....E  IPSRUNTM - Total Accumulated Run Time
X.....E  IPSSSCTM - CPU Use for System Calls
X.....E  IPSSYSTM - CPU Use in System Mode
X.....E  IPSTHCNT - Total Number of Thread Counts
X.....E  IPSUSRTM - CPU Use in User Mode
X.....E  IPSVSSKB - Virtual Memory Regions
```

```
+-----------------------------------------------------------------------------------------------------------------------------------+
|                                                                                                                                   |
| +-----------------+                                                                                                               |
| | I P S A L I T M |  Process Alive Time                                                                                           |
| +-----------------+                                                                                                               |
| IPSALITM contains the number of seconds the process was alive                                                                     |
| during the measurement interval.                                                                                                  |
|                                                                                                                                   |
| TYPE:  Accumulated - Numeric, measurement unit is SAS time                                                                        |
| value.                                                                                                                            |
|                                                                                                                                   |
| SEE ALSO:  Not applicable                                                                                                         |
|                                                                                                                                   |
| ELEMENT DERIVATION:  IPSALITM is set to the value of the                                                                          |
| MeasureWare Process metric PROC_INTERVAL_ALIVE.                                                                                   |
|                                                                                                                                   |
| RANGE OF VALUE:  IPSALITM is a SAS time value ranging from 0                                                                      |
| to a maximum that is limited only by practical bounds.                                                                            |
|                                                                                                                                   |
| SUMMARIZATION PROCESS:  IPSALITM contains the sum of all                                                                          |
| values of IPSALITM processed in the level of summarization.                                                                       |
|                                                                                                                                   |
| USAGE NOTES:  IPSALITM may be less than the time of the                                                                           |
| measurement interval if the process was new or died during                                                                        |
| that same interval.                                                                                                               |
|                                                                                                                                   |
| +-----------------+                                                                                                               |
| | I P S C P U T M |  Total CPU Use                                                                                                |
| +-----------------+                                                                                                               |
| IPSCPUTM contains the amount of CPU resources used by this                                                                        |
| process during the measurement interval.                                                                                          |
|                                                                                                                                   |
| TYPE:  Accumulated - Numeric, measurement unit is SAS time                                                                        |
| value.                                                                                                                            |
|                                                                                                                                   |
| SEE ALSO:  IPSUSRTM - CPU Use in User Mode                                                                                        |
|            IPSNICTM - CPU Use at Nice Priority                                                                                    |
|            IPSRTMTM - CPU Use at Real Time Priority                                                                               |
|            IPSCSWTM - CPU Use for Context Switching                                                                               |
|            IPSINTTM - CPU Use for Interrupt Handling                                                                              |
|            IPSSYSTM - CPU Use in System Mode                                                                                      |
|                                                                                                                                   |
| ELEMENT DERIVATION:  IPSCPUTM is set to the value of the                                                                          |
| MeasureWare Process metric PROC_CPU_TOTAL_TIME.                                                                                   |
|                                                                                                                                   |
| RANGE OF VALUE:  IPSCPUTM is a SAS time value ranging from 0                                                                      |
| to a maximum that is limited only by practical bounds.                                                                            |
|                                                                                                                                   |
| SUMMARIZATION PROCESS:  IPSCPUTM contains the sum of all                                                                          |
| values of IPSCPUTM processed in the level of summarization.                                                                       |
|                                                                                                                                   |
| USAGE NOTES:  None                                                                                                                |
|                                                                                                                                   |
| +-----------------+                                                                                                               |
| | I P S C S W T M |  CPU Use for Context Switching                                                                                |
| +-----------------+                                                                                                               |
| IPSCSWTM contains the number of seconds that the CPU was used                                                                     |
| for context switching during this process.                                                                                        |
|                                                                                                                                   |
| TYPE:  Accumulated - Numeric, measurement unit is SAS time                                                                        |
| value.                                                                                                                            |
|                                                                                                                                   |
| SEE ALSO:  IPSINTTM - CPU Use for Interrupt Handling                                                                              |
|            IPSNICTM - CPU Use at Nice Priorities                                                                                  |
|            IPSRTMTM - CPU Use at Real Time Priority                                                                               |
|            IPSUSRTM - CPU Use in User Mode                                                                                        |
|                                                                                                                                   |
| ELEMENT DERIVATION:  IPSCSWTM is set to the value of the                                                                          |
| MeasureWare Process metric PROC_CPU_CSWITCH_TIME.                                                                                 |
|                                                                                                                                   |
| RANGE OF VALUE:  IPSCSWTM is a SAS time value ranging from 0                                                                      |
| to a maximum that is limited by the number of seconds in                                                                          |
| the measurement interval.                                                                                                         |
|                                                                                                                                   |
| SUMMARIZATION PROCESS:  IPSCSWTM contains the sum of all                                                                          |
| values of IPSCSWTM processed in the level of summarization.                                                                       |
|                                                                                                                                   |
| USAGE NOTES:  None                                                                                                                |
|                                                                                                                                   |
| +-----------------+                                                                                                               |
| | I P S C U M T M |  Cumulative To-Date CPU Time                                                                                  |
| +-----------------+                                                                                                               |
| IPSCUMTM contains the total amount of CPU time used by the                                                                        |
| process to date.                                                                                                                  |
|                                                                                                                                   |
| TYPE:  Accumulated - Numeric, measurement unit is SAS time                                                                        |
| value.                                                                                                                            |
|                                                                                                                                   |
| SEE ALSO:  KPSRUNTM - Total Run Time                                                                                              |
|            KPSCPUTM - Total CPU Use                                                                                               |
|            IPSCPUTM - Total CPU Use                                                                                               |
|                                                                                                                                   |
| ELEMENT DERIVATION:  KPSCPUTM is set to the value of the                                                                          |
| MeasureWare Process metric PROC_CPU_TOTAL_TIME_CUM.                                                                               |
|                                                                                                                                   |
| RANGE OF VALUE:  IPSCUMTM is a SAS time value ranging from 0                                                                      |
| to a maximum that is limited only by practical bounds.                                                                            |
|                                                                                                                                   |
| SUMMARIZATION PROCESS:  IPSCUMTM contains the sum of all                                                                          |
| values of IPSCUMTM processed in the level of summarization.                                                                       |
|                                                                                                                                   |
|                                                                                                                                   |
+-----------------------------------------------------------------------------------------------------------------------------------+
```

USAGE NOTES: CA MICS separates processes into two files: the UNIX Process Activity (AUMKPS) file and the UNIX Process Exception (AUMIPS) file. The difference between these two files is that records for killed (terminated) processes are placed in the AUMKPS file; while records for processes that are still active, but are considered "interesting" by HP OpenView, are placed in the AUMIPS file.

The AUMIPS file is interval-based and many instances of a particular long-running process may appear in the AUMIPS file until final termination. Upon termination, a record for the long-running process appears in the AUMKPS file.

If the UNIX system crashes and final records are lost for the currently executing processes, the AUMIPS file can be searched for the last occurrence of an interval record for a particular, long-running process. The CPU time used over the lifetime of the process can be approximated by examining the value of IPSCUMTM. The value is approximate because it will not contain the CPU time consumed after the ENDTS of the AUMIPS interval record was found, up to the system crash.

The AUMKPS file, containing the records for killed processes, is suitable for billing purposes and overall process resource usage analysis. The KPSCPUTM data element in the AUMKPS file contains the total, cumulative CPU time used by the process.

The AUMIPS file, containing records for "interesting" processes during a particular measurement interval, also includes the data element, Total CPU Use (IPSCPUTM), which contains the CPU time used only during the interval.

IPSCUMTM is also included in the AUMIPS file to serve as a backup cumulative CPU time value for long-running processes in the event that a system crash or other problem prevents HP Openview from creating the final record for a process that is killed normally.

IPSCUMTM can only be active in the DETAIL timespan because it is a cumulative value. The value of IPSCUMTM for a particular process over several AUMIPS observations continually increases. The value in any particular AUMIPS observation represents the CPU time from the start of the process to the end of the interval represented.

| I P S D B K C T | Block I/Os

IPSDBKCT contains the number of block I/Os made by or for a process during the measurement interval. These are physical I/Os generated by file system access and do not include virtual memory I/O, or I/Os relating to raw disk access. These are I/Os for inode and superblock updates that are handled through the buffer cache.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDBRCT - Block Reads
IPSDBWCT - Block Writes

ELEMENT DERIVATION: IPSDBKCT is set to the value of the MeasureWare Process metric PROC_DISK_BLOCK_IO.

RANGE OF VALUE: IPSDBKCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSDBKCT contains the sum of all values of IPSDBKCT processed in the level of summarization.

USAGE NOTES: When a file is accessed on the AIX platform, it is memory mapped by the operating system. Accesses generate virtual memory I/Os. Reading a file generates block I/Os as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM I/Os) and posting updated inode information to disk (block I/Os).

| I P S D B R C T | Block Reads

IPSDBRCT contains the number of block reads made by a process during the measurement interval. These are physical reads generated by the file system and do not include virtual memory reads or raw reads. These are reads for inode and superblock data, handled through the buffer cache.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDBWCT - Block Writes

ELEMENT DERIVATION: IPSDBRCT is set to the value of the MeasureWare Process metric PROC_DISK_BLOCK_READ.

RANGE OF VALUE: IPSDBRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSDBRCT contains the sum of all values of IPSDBRCT processed in the level of summarization.

USAGE NOTES: When a file is accessed on the AIX platform, it is memory mapped by the operating system. Accesses generate virtual memory I/Os. Reading a file generates block I/Os as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM I/Os) and posting updated inode information to disk (block I/Os).

I P S D B W C T | Block Writes

IPSDBWCT contains the number of block writes made by the process during the interval. Calls destined for NFS mounted files are not included. These are physical writes generated by the file system and do not include virtual memory writes or writes relating to raw disk access. These are writes for inode and superblock updates, handled through the buffer cache.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDBRCT - Block Reads

ELEMENT DERIVATION: IPSDBWCT is set to the value of the MeasureWare Process metric PROC_DISK_BLOCK_WRITE.

RANGE OF VALUE: IPSDBWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSDBWCT contains the sum of all values of IPSDBWCT processed in the level of summarization.

USAGE NOTES: When a file is accessed, it is memory mapped by the operating system. Accesses generate virtual memory I/Os. Reading a file generates block I/Os as the file's inode information is cached. File writes are a combination of posting to memory mapped allocations (VM I/Os) and posting updated inode information to disk (block I/Os).

I P S D F R C T | File System Reads

IPSDFRCT contains the number of physical file reads issued by this process during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDFWCT - File System Writes
IPSLRDCT - Logical Disk Reads

ELEMENT DERIVATION: IPSDFRCT is set to the value of the MeasureWare Process metric PROC_DISK_FS_READ.

RANGE OF VALUE: IPSDFRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSDFRCT contains the sum of all values of IPSDFRCT processed in the level of summarization.

USAGE NOTES: The number of physical I/Os (IPSDFRCT) is usually smaller than the number of logical I/Os (IPSLRDCT).

I P S D F W C T | File System Writes

IPSDFWCT contains the number of physical file writes issued by this process during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDFRCT - File System Reads
IPSLWTCT - Logical Disk Writes

ELEMENT DERIVATION: IPSDFWCT is set to the value of the MeasureWare Process metric PROC_DISK_FS_READ.

RANGE OF VALUE: IPSDFWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSDFWCT contains the sum of all values of IPSDFWCT processed in the level of summarization.

USAGE NOTES: None

I P S D M M C T | Physical Memory Management I/Os

IPSDMMCT contains the number of physical memory management reads and writes that were issued by this process during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: IPSDMMCT is set to the value of the MeasureWare Process metric PROC_DISK_VM_IO.

RANGE OF VALUE: IPSDMMCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSDMMCT contains the sum of all values of IPSDMMCT processed in the level of summarization.

USAGE NOTES: None

I P S D S K I O | Physical Disk I/Os

IPSDSKIO contains the number of physical disk I/Os that occurred during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDMMCT - Physical Memory Management I/Os
IPSDFRCT - File System Reads
IPSDFWCT - File System Writes
IPSDSYCT - Physical File System Management I/Os

ELEMENT DERIVATION: IPSDSKIO is set to the value of the MeasureWare Process metric PROC_DISK_PHYS_IO.

RANGE OF VALUE: IPSDSKIO ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSDSKIO contains the sum of all values of IPSDSKIO processed in the level of summarization.

USAGE NOTES: None

I P S D S Y C T | Physical File System Management I/Os

IPSDSYCT contains the number of physical reads and writes for file system management that occured during this process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDMMCT - Physical Memory Management I/Os

ELEMENT DERIVATION: IPSDSYCT is set to the value of the MeasureWare Process metric PROC_DISK_SYSTEM_IO.

RANGE OF VALUE: IPSDSYCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSDSYCT contains the sum of all values of IPSDSYCT processed in the level of summarization.

USAGE NOTES: None

I P S F S D C T | File System Disk I/Os

IPSFSDCT contains the number of file system disk I/Os for this process during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDFRCT - File System Disk Reads
IPSDFWCT - File System Disk Writes

ELEMENT DERIVATION: IPSFSDCT is set to the value of the MeasureWare Process metric PROC_DISK_FS_IO.

RANGE OF VALUE: IPSFSDCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSFSDCT contains the sum of all values of IPSFSDCT processed in the level of summarization.

USAGE NOTES: None

I P S I N T R | Interest Reason Codes

IPSINTR contains the codes that explain why a process is interesting.

TYPE: Retained - Character, length 12.

SEE ALSO: IPSSTOPR - Blocked State

ELEMENT DERIVATION: IPSINTR is set to the value of the MeasureWare Process metric PROC_INTEREST.

RANGE OF VALUE: IPSINTR contains 12 positional values as defined in the MeasureWare metrics documentation, extracted below:

Place	Value	Meaning
1	N	New process
2	K	Killed process
3	C	CPU percentage used exceeded threshold
4	D	Disk I/Os exceeded threshold
5	P	Response-to-prompt time exceeded threshold
6	F	First-response time exceeded threshold
7	T	Transaction rate exceeded threshold
8	c	Wait for CPU percentage exceeded threshold
9	d	Wait for disk percentage exceeded threshold
10	m	Wait for memory percentage exceeded threshold
11	i	Wait for impede percentage exceeded threshold
12	blank	Reserved

Places 5-12 are set to a blank value and have no meaning for the SOL and AIX MeasureWare platforms.

SUMMARIZATION PROCESS: IPSINTR contains the last value of IPSINTR processed in the level of summarization.

USAGE NOTES: None

I P S I N T T M — CPU Use for Interrupt Handling

IPSINTTM contains the number of seconds that the CPU was used to handle interrupts while this process held the CPU during the measurement interval. The CPU time is charged to this process even if the process has nothing to do with the interrupt.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSCSWTM - CPU Use for Context Switching
IPSNICTM - CPU Use at Nice Priorities
IPSRTMTM - CPU Use at Real Time Priority
IPSUSRTM - CPU Use in User Mode

ELEMENT DERIVATION: IPSINTTM is set to the value of the MeasureWare Process metric PROC_CPU_INTERRUPT_TIME.

RANGE OF VALUE: IPSINTTM is a SAS time value ranging from 0 to a maximum that is limited by the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: IPSINTTM contains the sum of all values of IPSINTTM processed in the level of summarization.

USAGE NOTES: None

I P S L R D C T — Logical Disk Reads

IPSLRDCT contains the number of logical disk reads performed for this process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDFRCT - File System Reads
IPSLWTCT - Logical Disk Writes

ELEMENT DERIVATION: IPSLRDCT is set to the value of the MeasureWare Process metric PROC_DISK_LOGL_READ.

RANGE OF VALUE: IPSLRDCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSLRDCT contains the sum of all values of IPSLRDCT processed in the level of summarization.

USAGE NOTES: If buffered cache holds the data, physical I/O will not occur, so a logical read does not always result in a physical I/O.

I P S L W T C T — Logical Disk Writes

IPSLWTCT contains the number of logical disk writes performed for this process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDFRCT - File System Reads
IPSDFWCT - File System Writes
IPSLRDCT - Logical Disk Reads

ELEMENT DERIVATION: IPSLWTCT is set to the value of the MeasureWare Process metric PROC_DISK_LOGL_WRITE.

RANGE OF VALUE: IPSLWTCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSLWTCT contains the sum of all values of IPSLWTCT processed in the level of summarization.

USAGE NOTES: If the amount of data being logically written is smaller than the size of the file system block and the block to which the data is being written is not in the buffer cache, the logical write causes a physical read to fetch the block and a physical write to write the block on disk.

To maximize efficiency, write blocks that are the same size as the file system blocks.

I P S N I C T M | CPU Use at Nice Priority

IPSNICTM contains the number of seconds that the CPU was used by this process at a nice priority.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSCSWTM - CPU Use for Context Switching
IPSINTTM - CPU Use for Interrupt Handling
IPSRTMTM - CPU Use at Real Time Priority
IPSUSRTM - CPU Use in User Mode

ELEMENT DERIVATION: IPSNICTM is set to the value of the MeasureWare Process metric PROC_CPU_NICE_TIME.

RANGE OF VALUE: IPSNICTM is a SAS time value ranging from 0 to a maximum that is limited to the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: IPSNICTM contains the sum of all values of IPSNICTM processed in the level of summarization.

USAGE NOTES: None

I P S N O R T M | CPU Use at Normal Priority

IPSNORTM contains the number of seconds that the CPU was used to run programs at normal priority for this process during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSNICTM - CPU Use at Nice Priority
IPSRTMTM - CPU Use at Real Time Priority

ELEMENT DERIVATION: IPSNORTM is set to the value of the MeasureWare Process metric PROC_CPU_NORMAL_TIME.

RANGE OF VALUE: IPSNORTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSNORTM contains the sum of all values of IPSNORTM processed in the level of summarization.

USAGE NOTES: None

I P S P F M A J | Page Faults to Disk

IPSPFMAJ contains the number of page faults to disk that occurred during this process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSPFMIN - Page Faults to Memory

ELEMENT DERIVATION: IPSPFMAJ is set to the value of the MeasureWare Process metric PROC_MAJOR_FAULT.

RANGE OF VALUE: IPSPFMAJ takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSPFMAJ contains the sum of all values of IPSPFMAJ processed in the level of summarization.

USAGE NOTES: Page Faults to Disk result in at least one physical disk I/O. Because more than one page can be read when a major fault occurs, diskless faults result the first time the extra pages are read.

I P S P F M I N | Page Faults to Memory

IPSPFMIN contains the number of page faults to memory that occured during this process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSPFMAJ - Page Faults to Disk

ELEMENT DERIVATION: IPSPFMIN is set to the value of the MeasureWare Process metric PROC_MINOR_FAULT.

RANGE OF VALUE: IPSPFMIN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSPFMIN contains the sum of all values of IPSPFMIN processed in the level of summarization.

USAGE NOTES: Page faults to memory do not result in disk I/O.

I P S P K B C T | I/O Transfer Count in KBytes

IPSPKBCT contains the number of kilobytes of data read or written by a process during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSDSKIO - Physical Disk I/Os

ELEMENT DERIVATION: IPSPKBCT is set to the value of the MeasureWare Process metric PROC_IO_BYTE.

RANGE OF VALUE: IPSPKBCT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSPKBCT contains the sum of all values of IPSPKBCT processed in the level of summarization.

USAGE NOTES: None

I P S P R M I D | PRM Group Identification Number

IPSPRMID is specific to a process and contains the Process Resource Manager (PRM) group identifier.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: IPSPRMID is set to the value of the MeasureWare Process metric PROC_PRMID.

RANGE OF VALUE: IPSPRMID can contain any combination of 1 to 24 alphanumeric characters.

SUMMARIZATION PROCESS: IPSPRMID contains the last value of IPSPRMID processed in the level of summarization.

USAGE NOTES: The PRM group configuration is kept in the PRM configuration file.

I P S R S S K B | Resident Set Size

IPSRSSKB contains the size (in kilobytes) of the process's data, stack, and text (code) pages that resided in memory at the end of the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: IPSVSSKB - Virtual Memory Regions

ELEMENT DERIVATION: IPSRSSKB is set to the value of the MeasureWare Process metric PROC_MEM_RES.

RANGE OF VALUE: IPSRSSKB takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSRSSKB contains the sum of all values of IPSRSSKB processed in the level of summarization.

USAGE NOTES: None

I P S R T M T M | CPU Use at Real Time Priority

IPSRTMTM contains the number of seconds that the CPU spent running this process at real time priorities.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSCSWTM - CPU Use for Context Switching
IPSINTTM - CPU Use for Interrupt Handling
IPSNICTM - CPU Use at Nice Priorities
IPSUSRTM - CPU Use in User Mode

ELEMENT DERIVATION: IPSRTMTM is set to the value of the

MeasureWare Process metric PROC_CPU_REALTIME_TIME.

RANGE OF VALUE: IPSRTMTM is a SAS time value ranging from 0 to a maximum that is limited by the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: IPSRTMTM contains the sum of all values of IPSRTMTM processed in the level of summarization.

USAGE NOTES: None

I P S R U N T M | Total Accumulated Run Time

IPSRUNTM contains the total number of seconds the process has been running (wall clock time).

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: IPSRUNTM is set to the value of the MeasureWare Process metric PROC_RUN_TIME.

RANGE OF VALUE: IPSRUNTM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSRUNTM contains the sum of all values of IPSRUNTM processed in the level of summarization.

USAGE NOTES: None

I P S S S C T M | CPU Use for System Calls

IPSSSCTM contains the number of seconds that the CPU was in system mode for this process during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSUSRTM - CPU Use in User Mode
IPSNORTM - CPU Use at Normal Priority
IPSNICTM - CPU Use at Nice Priority
IPSRTMTM - CPU Use at Real Time Priority
IPSSYSTM - CPU Use in System Mode
IPSCSWTM - CPU Use for Context Switching
IPSINTTM - CPU Use for Interrupt Handling
IPSCPUTM - Total CPU Use

ELEMENT DERIVATION: IPSSSCTM is set to the value of the MeasureWare Process metric PROC_CPU_SYSCALL_TIME.

RANGE OF VALUE: IPSSSCTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSSSCTM contains the sum of all values of IPSSSCTM processed in the level of summarization.

USAGE NOTES: None

I P S S T O P R | Blocked State

IPSSTOPR contains the codes that explain why a process lost the CPU entering a blocked state the last time during the measurement interval.

TYPE: Retained - Character, length 12.

SEE ALSO: IPSINTR - Interest Reason Codes

ELEMENT DERIVATION: IPSSTOPR is set to the value of the MeasureWare Process metric PROC_STOP_REASON.

RANGE OF VALUE: IPSSTOPR can contain any combination of 1 to 12 alphanumeric characters as defined in the MeasureWare metrics documentation.

A complete list of block reasons on a HP-UX platform are listed below:

Cache	Cache buffer availability
CDFS	CD-ROM file system node structure allocation or locks
Died	Process terminated during the interval
Disk	Wait for an I/O operation to complete
Graph	Wait for graphics card or framebuf semaphore to complete
Inode	Wait for an Inode operation to complete
IO	Wait for non-disk I/O to complete
IPC	Wait for interprocess or remote process operation to complete
JOBCL	Wait for debug resume or job control start
LAN	Wait for a LAN operation to complete
Mbuf	Wait for memory buffer
Mesg	Wait for a Message queue operation

New	Process created during the interval
NFS	Wait for a Network File System request
None	Process is ready to run
Other	Unknown because process started before midaemon
Pipe	Wait for data from a pipe
Pri	Wait for higher priority process
RPC	Wait for remote procedure call operation
Sem	Wait for a semaphore operation
Sleep	Wait for an event to complete
Sockt	Wait for a socket operation to complete
Strms	Wait for operation to complete while accessing a "streams" device
Systm	Wait for system resources or locks
Term	Wait for a terminal transfer
VM	Wait for a Virtual memory operation

A complete list of block reasons on a SunOS platform are listed below:

Died	Process terminated during an interval
New	Process created during the interval
None	Process is ready to run
Other	Unknown because process started before midaemon
Pmem	Waiting for additional primary memory
Pri	Process is on the run queue
Sleep	Wait for an event to complete
Trace	Received a signal to stop because of ongoing trace
Zomb	Process has terminated

SUMMARIZATION PROCESS: IPSSTOPR contains the last value of IPSSTOPR processed in the level of summarization.

USAGE NOTES: None

I P S S Y S T M — CPU Use in System Mode

IPSSYSTM contains the number of seconds that the CPU was used in system (kernel) mode during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSSCPTM - CPU Wait Time

ELEMENT DERIVATION: IPSSYSTM is set to the value of the MeasureWare Process metric PROC_CPU_SYS_MODE_TIME.

RANGE OF VALUE: IPSSYSTM is a SAS time value ranging from 0 to a maximum that is limited by the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: IPSSYSTM contains the sum of all values of IPSSYSTM processed in the level of summarization.

USAGE NOTES: None

I P S T H C N T — Total Number of Thread Counts

IPSTHCNT contains the total number of kernel threads for the current process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: IPSTHCNT is set to the value of the MeasureWare Process metric PROC_THREAD_COUNT.

RANGE OF VALUE: IPSTHCNT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSTHCNT contains the sum of all values of IPSTHCNT processed in the level of summarization.

USAGE NOTES: None

I P S U S R T M — CPU Use in User Mode

IPSUSRTM contains the number of seconds that the CPU was used in user mode for this process during this measurement interval.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: IPSCSWTM - CPU Use for Context Switching
IPSINTTM - CPU Use for Interrupt Handling
IPSNICTM - CPU Use at Nice Priorities
IPSRTMTM - CPU Use at Real Time Priority

ELEMENT DERIVATION: IPSUSRTM is set to the value of the MeasureWare Process metric PROC_CPU_USER_MODE_TIME.

RANGE OF VALUE: IPSUSRTM is a SAS time value ranging from 0

to a maximum that is limited by the number of seconds in the measurement interval.

SUMMARIZATION PROCESS: IPSUSRTM contains the sum of all values of IPSUSRTM processed in the level of summarization.

USAGE NOTES: None

I P S V S S K B | Virtual Memory Regions

IPSVSSKB contains the number of kilobytes of virtual memory reserved for this process.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: IPSVSSKB is set to the value of the MeasureWare Process metric PROC_MEM_VIRT.

RANGE OF VALUE: IPSVSSKB takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: IPSVSSKB contains the sum of all values of IPSVSSKB processed in the level of summarization.

USAGE NOTES: A value of "na" may be shown for the SunOS platform when this information is unobtainable, or for the HP-UX platform when a swapper process is run.

B.6 - The UNIX LAN Activity (AUMNIF) File

The UNIX LAN Activity file contains interval-based measurements useful for capacity analysis. It includes information about LAN packets, errors, and collisions, and is derived from MeasureWare NETIF metrics.

The metrics supported vary by operating system.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	NIFNAME	YEAR	MONTH
	DAY	HOUR	ENDTS		
DAYS	SYSID	AUMOPSYS	NIFNAME	YEAR	MONTH
	DAY	HOUR			
WEEKS	SYSID	AUMOPSYS	NIFNAME	YEAR	WEEK
	ZONE	HOUR			
MONTHS	SYSID	AUMOPSYS	NIFNAME	YEAR	MONTH
	ZONE				
YEARS	SYSID	AUMOPSYS	NIFNAME	YEAR	ZONE
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	NIFNAME	- Device Name for LAN Interface
XDWMY.E	SYSID	- System Identifier
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
		Common Data Elements
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
		Retained Data Elements
XDWMY.E	NIFMTUSZ	- MTU Size
XDWMY.E	NIFQUEUE	- Outbound Queue Length
XDWMY.E	NIFSPEED	- Interface Speed
		Accumulated Data Elements
XDWMY.E	NIFCOLCT	- LAN Collisions
XDWMY.E	NIFERRCT	- LAN Errors
XDWMY.E	NIFPKICT	- Inbound LAN Packets
XDWMY.E	NIFPKOCT	- Outbound LAN Packets
		Derived Data Elements
XDWMY.E	NIFPSCOL	- LAN Collision Rate
XDWMY.E	NIFPSERR	- LAN Error Rate
XDWMY.E	NIFPSPKI	- Inbound LAN Packet Rate
XDWMY.E	NIFPSPKO	- Outbound LAN Packet Rate

N I F C O L C T — LAN Collisions

NIFCOLCT contains the number of LAN collisions that occurred during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: NIFERRCT - LAN Errors

ELEMENT DERIVATION: NIFCOLCT is set to the value of the MeasureWare NETIF metric BYNETIF_COLLISION.

RANGE OF VALUE: NIFCOLCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: NIFCOLCT contains the sum of all values of NIFCOLCT processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| N I F E R R C T |  LAN Errors
+-----------------+
NIFERRCT contains the number of LAN errors that occurred
during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  NIFCOLCT - LAN Collisions

ELEMENT DERIVATION:  NIFERRCT is set to the value of the
MeasureWare NETIF metric BYNETIF_ERROR.

RANGE OF VALUE:  NIFERRCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  NIFERRCT contains the sum of all
values of NIFERRCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| N I F M T U S Z |  MTU Size
+-----------------+
NIFMTUSZ contains the Maximum Transfer Unit (MTU) size.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  NIFMTUSZ is set to the value of the
MeasureWare NETIF metric BYNETIF_NET_MTU.

RANGE OF VALUE:  NIFMTUSZ ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  NIFMTUSZ contains the last value of
NIFMTUSZ processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| N I F N A M E   |  Device Name for LAN Interface
+-----------------+
NIFNAME contains the device name of the LAN interface.

TYPE:  Retained - Character, length 20.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  NIFNAME is set to the value of the
MeasureWare NETIF metric BYNETIF_NAME.

RANGE OF VALUE:  NIFNAME can contain any combination of 1 to
20 characters.

SUMMARIZATION PROCESS:  NIFNAME contains the last value of
NIFNAME processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| N I F P K I C T |  Inbound LAN Packets
+-----------------+
NIFPKICT contains the number of inbound LAN packets sent to
this LAN.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  NIFPKOCT - Outbound LAN Packets

ELEMENT DERIVATION:  NIFPKICT is set to the value of the
MeasureWare NETIF metric BYNETIF_IN_PACKET.

RANGE OF VALUE:  NIFPKICT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  NIFPKICT contains the sum of all
values of NIFPKICT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| N I F P K O C T |  Outbound LAN Packets
+-----------------+
NIFPKOCT contains the number of outbound LAN packets sent
from this LAN.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  NIFPKICT - Inbound LAN Packets

ELEMENT DERIVATION:  NIFPKOCT is set to the value of the
MeasureWare NETIF metric BYNETIF_OUT_PACKET.

RANGE OF VALUE:  NIFPKOCT takes any numeric value ranging
```

from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: NIFPKOCT contains the sum of all values of NIFPKOCT processed in the level of summarization.

USAGE NOTES: None

N I F P S C O L | LAN Collision Rate

NIFPSCOL contains the rate at which LAN collisions occured.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time
NIFCOLCT - LAN Collisions

ELEMENT DERIVATION: NIFPSCOL is computed using the following algorithm:

```
NIFPSCOL = DURATION;
IF DURATION
  THEN NIFPSCOL = NIFCOLCT/DURATION;
```

RANGE OF VALUE: NIFPSCOL takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: NIFPSCOL contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

N I F P S E R R | LAN Error Rate

NIFPSERR contains the rate at which LAN errors occurred on this LAN.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time
NIFERRCT - LAN Errors

ELEMENT DERIVATION: NIFPSERR is computed using the following algorithm:

```
NIFPSERR = DURATION;
IF DURATION
  THEN NIFPSERR = NIFERRCT/DURATION;
```

RANGE OF VALUE: NIFPSERR takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: NIFPSERR contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

N I F P S P K I | Inbound LAN Packet Rate

NIFPSPKI contains the rate at which LAN packets were sent to this LAN.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time
NIFPKICT - Inbound LAN Packets

ELEMENT DERIVATION: NIFPSPKI is computed using the following algorithm:

```
NIFPSPKI = DURATION;
IF DURATION
  THEN NIFPSPKI = NIFPKICT/DURATION;
```

RANGE OF VALUE: NIFPSPKI takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: NIFPSPKI contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

N I F P S P K O | Outbound LAN Packet Rate

NIFPSPKO contains the rate at which packets were sent from this LAN.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DURATION - Recording Interval Time
NIFPKOCT - Outbound LAN Packets

ELEMENT DERIVATION: NIFPSPKO is computed using the following algorithm:

```
NIFPSPKO = DURATION;
IF DURATION
  THEN NIFPSPKO = NIFPKOCT/DURATION;
```

RANGE OF VALUE: NIFPSPKO takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: NIFPSPKO contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

N I F Q U E U E — Outbound Queue Length

NIFQUEUE contains the outbound queue length at the time of the last sample, regardless of the number of IP addresses on the system.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: NIFQUEUE is set to the value of the MeasureWare NETIF metric BYNETIF_QUEUE.

RANGE OF VALUE: NIFQUEUE ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: NIFQUEUE contains the last value of NIFQUEUE processed in the level of summarization.

USAGE NOTES: None

N I F S P E E D — Interface Speed

NIFSPEED contains the number of physical collisions that occurred on the network interface during the measurement interval.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: NIFSPEED is set to the value of the MeasureWare NETIF metric BYNETIF_COLLISION.

RANGE OF VALUE: NIFSPEED ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: NIFSPEED contains the last value of NIFSPEED processed in the level of summarization.

USAGE NOTES: NIFSPEED does not currently include deferred packets.

B.7 - The UNIX Disk Activity (AUMDDA) File

The UNIX Disk Activity file contains global system measurements useful for capacity analysis. It includes information about logical and physical reads and writes, as well as utilization rates. It is derived from MeasureWare Global metrics.

The metrics supported vary by operating system.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	DDANAME	YEAR	MONTH
	DAY	HOUR	ENDTS		
DAYS	SYSID	AUMOPSYS	DDANAME	YEAR	MONTH
	DAY	HOUR			
WEEKS	SYSID	AUMOPSYS	DDANAME	YEAR	WEEK
	ZONE	HOUR			
MONTHS	SYSID	AUMOPSYS	DDANAME	YEAR	MONTH
	ZONE				
YEARS	SYSID	AUMOPSYS	DDANAME	YEAR	ZONE
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDWMY.E	DDANAME	- Disk Name
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	SYSID	- System Identifier
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
		Common Data Elements
XDWMY.E	AUMLDEV	- Logical Device Number
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
		Retained Data Elements
XDWMY.E	DDADIR	- File System Directory
		Accumulated Data Elements
XDWMY.E	DDADDPDR	- Time Processing Disk Request
XDWMY.E	DDADFRCT	- File System Reads
XDWMY.E	DDADFWCT	- File System Writes
XDWMY.E	DDADMMCT	- Virtual Memory I/Os
XDWMY.E	DDADRRCT	- Raw Reads
XDWMY.E	DDADRWCT	- Raw Writes
XDWMY.E	DDADSKIO	- Physical I/Os
XDWMY.E	DDADUTTM	- Disk Busy Time
XDWMY.E	DDALRDCT	- Logical Disk Reads
XDWMY.E	DDALWTCT	- Logical Disk Writes
XDWMY.E	DDAPKBCT	- Disk I/O Transfer Count in KBytes
XDWMY.E	DDAPRDCT	- Physical Reads
XDWMY.E	DDAPWTCT	- Physical Writes
XDWMY.E	DDAQLNSM	- Sum of Avg Queue Length
XDWMY.E	DDARKBCT	- KBytes Read
XDWMY.E	DDASYSCT	- File System Management I/Os
XDWMY.E	DDAWKBCT	- KBytes Written
		Derived Data Elements
XDWMY.E	DDAAVQLN	- Avg Queue Length
XDWMY.E	DDADSKRT	- Physical I/O Rate
XDWMY.E	DDAPCDUT	- Percent Disk Utilization
XDWMY.E	DDAPKBRT	- Disk I/O Transfer Rate in KBytes
XDWMY.E	DDAPSFRD	- File System Read Rate
XDWMY.E	DDAPSFWT	- File System Write Rate
XDWMY.E	DDAPSLRD	- Logical Disk Read Rate
XDWMY.E	DDAPSLWT	- Logical Disk Write Rate
XDWMY.E	DDAPSMEM	- Virtual Memory I/O Rate
XDWMY.E	DDAPSRRD	- Raw Read Rate
XDWMY.E	DDAPSRWT	- Raw Write Rate
XDWMY.E	DDAPSSYS	- File System Management I/O Rate
XDWMY.E	DDARKBRT	- Disk Read Transfer Rate in KBytes
XDWMY.E	DDAWKBRT	- Disk Write Transfer Rate in KBytes

```
+-----------------+
| D D A A V Q L N |  Avg Queue Length
+-----------------+
DDAAVQLN contains the average length of the queue waiting for
access to this disk.  This measures processes that are
blocked waiting for physical I/O to complete using the disk
blocked time as an average of the total execution for all
processes.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  INTERVLS - Number of Recording Intervals
           DDAQLNSM - Sum of Avg Queue Length

ELEMENT DERIVATION:  DDAAVQLN is computed using the following
algorithm:

  DDAAVQLN = INTERVLS;
  IF INTERVLS
    THEN DDAAVQLN = DDAQLNSM/INTERVLS;

RANGE OF VALUE:  DDAAVQLN takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAAVQLN contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| D D A D D P D R |  Time Processing Disk Request
+-----------------+
DDADDPDR contains the time spent by the disk device
processing each disk request during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is SAS time
value.

SEE ALSO:  DURATION - Recording Interval Time
           GPRAVPRQ - Avg Processes Blocked on Priority

ELEMENT DERIVATION:  At the DETAIL timespan, DDADDPDR is
computed using MeasureWare Disk metrics in the following
calculation:

  BYDSK_AVG_SERVICE_TIME*DURATION

RANGE OF VALUE:  DDADDPDR is a SAS time value ranging from 0
to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDADDPDR contains the sum of all
values of DDADDPDR processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| D D A D F R C T |  File System Reads
+-----------------+
DDADFRCT contains the number of times this disk was read by
the file system.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  DDADRRCT - Raw Reads
           DDALRDCT - Logical Disk Reads

ELEMENT DERIVATION:  DDADFRCT is set to the value of the
MeasureWare Disk metric BYDSK_FS_READ.

RANGE OF VALUE:  DDADFRCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDADFRCT contains the sum of all
values of DDADFRCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| D D A D F W C T |  File System Writes
+-----------------+
DDADFWCT contains the number of times the file system wrote
to this disk.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  DDALWTCT - Logical Disk Writes
           DDADRWCT - Raw Writes

ELEMENT DERIVATION:  DDADFWCT is set to the value of the
MeasureWare Disk metric BYDSK_FS_WRITE.

RANGE OF VALUE:  DDADFWCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDADFWCT contains the sum of all
values of DDADFWCT processed in the level of summarization.
```

USAGE NOTES: None

| D D A D I R | File System Directory

DDADIR contains the name of the file system directory mounted on this device.

TYPE: Retained - Character, length 40.

SEE ALSO: Not applicable

ELEMENT DERIVATION: DDADIR is set to the value of the MeasureWare Disk metric BYDSK_DIRNAME.

RANGE OF VALUE: DDADIR can contain any combination of 1 to 40 alphanumeric characters.

SUMMARIZATION PROCESS: DDADIR contains the last value of DDADIR processed in the level of summarization.

USAGE NOTES: If more than one file system is mounted on this device, "multiple FS" is seen.

| D D A D M M C T | Memory Manager Transfers

DDADMMCT contains the number of physical disk reads and physical disk writes for memory management tasks.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: DDASYSCT - System Transfers

ELEMENT DERIVATION: DDADMMCT is set to the value of the MeasureWare Disk metric BYDSK_VM_IO.

RANGE OF VALUE: DDADMMCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDADMMCT contains the sum of all values of DDADMMCT processed in the level of summarization.

USAGE NOTES: None

| D D A D R R C T | Raw Reads

DDADRRCT contains the number of raw reads to this disk.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: DDALRDCT - Logical Disk Reads
DDADFRCT - File System Reads

ELEMENT DERIVATION: DDADRRCT is set to the value of the MeasureWare Disk metric BYDSK_RAW_READ.

RANGE OF VALUE: DDADRRCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDADRRCT contains the sum of all values of DDADRRCT processed in the level of summarization.

USAGE NOTES: None

| D D A D R W C T | Raw Writes

DDADRWCT contains the number of raw writes to this disk.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: DDALWTCT - Logical Disk Writes
DDADFWCT - File System Writes

ELEMENT DERIVATION: DDADRWCT is set to the value of the MeasureWare Disk metric BYDSK_RAW_WRITE.

RANGE OF VALUE: DDADRWCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDADRWCT contains the sum of all values of DDADRWCT processed in the level of summarization.

USAGE NOTES: None

| D D A D S K I O | Number of Physical I/Os

DDADSKIO contains the number of physical I/Os for this disk device during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not available

ELEMENT DERIVATION: DDADSKIO is set to the value of the MeasureWare Disk metric BYDSK_PHYS_IO.

RANGE OF VALUE: DDADSKIO ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDADSKIO contains the sum of all values of DDADSKIO processed in the level of summarization.

USAGE NOTES: None

D D A D S K R T — Physical I/O Rate

DDADSKRT contains the physical I/O rate for this disk during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is integer.

SEE ALSO: DDADSKIO - Number of Physical I/Os
DURATION - Recording Interval Time

ELEMENT DERIVATION: DDADSKRT is computed using the following algorithm:

```
DDADSKRT = DURATION;
IF DURATION
  THEN DDADSKRT = DDADSKIO/DURATION;
```

RANGE OF VALUE: DDADSKRT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDADSKRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

D D A D U T T M — Disk Busy Time

DDADUTTM contains the amount of time this disk was busy.

TYPE: Accumulated - Numeric, measurement unit is SAS time value.

SEE ALSO: DDAPCDUT - Percent Disk Utilization
DURATION - Recording Interval Time

ELEMENT DERIVATION: At the DETAIL time-span, DDADUTTM is computed using MeasureWare metrics in the following calculation:

```
(BYDSK_UTIL*INTERVAL)/100
```

where BYDSK_UTIL is from Disk metrics and INTERVAL is from Global metrics.

RANGE OF VALUE: DDADUTTM is a SAS time value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDADUTTM contains the sum of all values of DDADUTTM processed in the level of summarization.

USAGE NOTES: None

D D A L R D C T — Logical Disk Reads

DDALRDCT contains the number of logical reads to this disk.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: DDALWTCT - Logical Disk Writes
DDADRRCT - Raw Reads

ELEMENT DERIVATION: DDALRDCT is set to the value of the MeasureWare Disk metric BYDSK_LOGL_READ.

RANGE OF VALUE: DDALRDCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDALRDCT contains the sum of all values of DDALRDCT processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| D D A L W T C T |  Logical Disk Writes
+-----------------+
DDALWTCT contains the number of logical disk writes to this
disk.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  DDALRDCT - Logical Disk Reads
           DDADRWCT - Raw Writes

ELEMENT DERIVATION:  DDALWTCT is set to the value of the
MeasureWare Disk metric BYDSK_LOGL_WRITE.

RANGE OF VALUE:  DDALWTCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDALWTCT contains the sum of all
values of DDALWTCT processed in the level of summarization.

USAGE NOTES:  None



+-----------------+
| D D A N A M E   |  Disk Name
+-----------------+
DDANAME contains the interface type and logical unit number
of the disk device.

TYPE:  Retained - Character, length 40.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  DDANAME is set to the value of the
MeasureWare Disk metric BYDSK_DEVNAME.

RANGE OF VALUE:  DDANAME can contain any combination of 1 to
40 alphanumeric characters.

SUMMARIZATION PROCESS:  DDANAME contains the last value of
DDANAME processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| D D A P C D U T |  Percent Disk Utilization
+-----------------+
DDAPCDUT contains percentage of time this disk was used
during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DDADUTTM - Disk Busy Time
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  DDAPCDUT is computed using the following
algorithm:

  DDAPCDUT = DURATION;
  IF DURATION
    THEN DDAPCDUT = 100*DDADUTTM/DURATION;

RANGE OF VALUE:  DDAPCDUT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAPCDUT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None



+-----------------+
| D D A P K B C T |  Disk IO Transfer Count in KBytes
+-----------------+
DDAPKBCT contains the average KBytes transferred to or from
the current disk device during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  DDARKBCT - KBytes Read
           DDAPKBRT - Disk I/O Transfer Rate in KBytes
           DDAWKBCT - KBytes Written

ELEMENT DERIVATION:  DDAPKBCT is set to the value of the
MeasureWare Application metric BYDSK_PHYS_BYTE.

RANGE OF VALUE:  DDAPKBCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAPKBCT contains the sum of all
values of DDAPKBCT processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| D D A P K B R T |  Disk I/O Transfer Rate in KBytes
+-----------------+
DDAPKBRT contains the average KBytes per second transferred
to or from the current disk device during the measurement
interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DDAPKBCT - Disk I/O Transfer Count in KBytes
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  DDAPKBRT is computed using the following
algorithm:

  DDAPKBRT = DURATION;
  IF DURATION
    THEN DDAPKBRT = DDAPKBCT/DURATION;

RANGE OF VALUE:  DDAPKBRT ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAPKBRT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| D D A P R D C T |  Physical Reads
+-----------------+
DDAPRDCT contains the number of physical reads for the disk
device during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  DDAPWTCT - Physical Writes
           DDADSKIO - Physical I/Os

ELEMENT DERIVATION:  DDAPRDCT is set to the value of the
MeasureWare metric BYDSK_PHYS_READ.

RANGE OF VALUE:  DDAPRDCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAPRDCT contains the sum of all
values of DDAPRDCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| D D A P S F R D |  File System Read Rate
+-----------------+
DDAPSFRD contains the rate at which the file system read data
from this disk.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DDADFRCT - File System Reads
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  DDAPSFRD is computed using the following
algorithm:

  DDAPSFRD = DURATION;
  IF DURATION
    THEN DDAPSFRD = DDADFRCT/DURATION;

RANGE OF VALUE:  DDAPSFRD takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAPSFRD contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| D D A P S F W T |  File System Write Rate
+-----------------+
DDAPSFWT contains the rate at which the file system wrote to
this disk.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DDADFWCT - File System Writes
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  DDAPSFWT is computed using the following
algorithm:

  DDAPSFWT = DURATION;
```

```
  IF DURATION
    THEN DDAPSFWT = DDADFWCT/DURATION;
```

RANGE OF VALUE: DDAPSFWT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDAPSFWT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

D D A P S L R D — Logical Disk Read Rate

DDAPSLRD contains the rate at which this disk was logically read during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DDALRDCT - Logical Disk Reads
DURATION - Recording Interval Time

ELEMENT DERIVATION: DDAPSLRD is computed using the following algorithm:

```
  DDAPSLRD = DURATION;
  IF DURATION
    THEN DDAPSLRD = DDALRDCT/DURATION;
```

RANGE OF VALUE: DDAPSLRD takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDAPSLRD contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

D D A P S L W T — Logical Disk Write Rate

DDAPSLWT contains the rate to which this disk was logically written during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DDALWTCT - Logical Disk Writes
DURATION - Recording Interval Time

ELEMENT DERIVATION: DDAPSLWT is computed using the following algorithm:

```
  DDAPSLWT = DURATION;
  IF DURATION
    THEN DDAPSLWT = DDALWTCT/DURATION;
```

RANGE OF VALUE: DDAPSLWT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDAPSLWT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

D D A P S M E M — Memory Manager Transfer Rate

DDAPSMEM contains the rate at which physical disk transfers for memory management tasks occurred.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DDADMMCT - Memory Manager Transfers
DURATION - Recording Interval Time

ELEMENT DERIVATION: DDAPSMEM is computed using the following algorithm:

```
  DDAPSMEM = DURATION;
  IF DURATION
    THEN DDAPSMEM = DDADMMCT/DURATION;
```

RANGE OF VALUE: DDAPSMEM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDAPSMEM contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

```
+-----------------+
| D D A P S R R D |  Raw Read Rate
+-----------------+
DDAPSRRD contains the rate at which raw data was read from
this disk.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DDADRRCT - Raw Reads
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  DDAPSRRD is computed using the following
algorithm:

  DDAPSRRD = DURATION;
  IF DURATION
    THEN DDAPSRRD = DDADRRCT/DURATION;

RANGE OF VALUE:  DDAPSRRD takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAPSRRD contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None



+-----------------+
| D D A P S R W T |  Raw Write Rate
+-----------------+
DDAPSRWT contains the rate at which raw data was written to
this disk.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DDADRWCT - Raw Writes
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  DDAPSRRD is computed using the following
algorithm:

  DDAPSRRD = DURATION;
  IF DURATION
    THEN DDAPSRRD = DDADRWCT/DURATION;

RANGE OF VALUE:  DDAPSRRD takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.
```

```
SUMMARIZATION PROCESS:  DDAPSRRD contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None



+-----------------+
| D D A P S S Y S |  System Transfer Rate
+-----------------+
DDAPSSYS contains the rate at which the system transferred
data to and from this disk.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DDASYSCT - System Transfers
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  DDAPSSYS is computed using the following
algorithm:

  DDAPSSYS = DURATION;
  IF DURATION
    THEN DDAPSSYS = DDASYSCT/DURATION;

RANGE OF VALUE:  DDAPSSYS takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAPSSYS contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None



+-----------------+
| D D A P W T C T |  Physical Writes
+-----------------+
DDAPWTCT contains the number of physical writes for this disk
device during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is
integer.

SEE ALSO:  DDAPRDCT - Physical Reads
           DDADSKIO - Physical I/Os

ELEMENT DERIVATION:  DDAPWTCT is set to the value of the
MeasureWare metric BYDSK_PHYS_WRITE.
```

RANGE OF VALUE: DDAPWTCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDAPWTCT contains the sum of all values of DDAPWTCT processed in the level of summarization.

USAGE NOTES: None

D D A Q L N S M — Sum of Avg Queue Length

DDAQLNSM contains the sum of the average queue length for this disk that was sampled during this measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: AUMSAMPL - Number of Data Samples
DDAAVQLN - Average Queue Length

ELEMENT DERIVATION: At the DETAIL time-span, DDAQLNSM is computed using MeasureWare metrics in the following calculation:

```
BYDSK_REQUEST_QUEUE*GBL_PROC_SAMPLE
```

where BYDSK_REQUEST_QUEUE is from Disk data and GBL_PROC_SAMPLE is from Global data.

RANGE OF VALUE: DDAQLNSM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDAQLNSM contains the sum of all values of DDAQLNSM processed in the level of summarization.

USAGE NOTES: None

D D A R K B C T — KBytes Read

DDARKBCT contains the number of KBytes transferred from the disk device during the measurement interval. This counts all types of disk reads, including file system, virtual memory, and raw IO.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: DDAWKBCT - KBytes Written
DDARKBRT - Disk Read Transfer Rate in KBytes

ELEMENT DERIVATION: DDARKBCT is set to the value of the MeasureWare Application metric BYDSK_PHYS_READ_BYTE.

RANGE OF VALUE: DDARKBCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDARKBCT contains the sum of all values of DDARKBCT processed in the level of summarization.

USAGE NOTES: None

D D A R K B R T — Disk Read Transfer Rate in KBytes

DDARKBRT contains the average KBytes per second transferred to the disk during the measurement interval. This count includes all types of writes, including file system, virtual memory, and raw IO.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: DDARKBCT - KBytes Read
DURATION - Recording Interval Time

ELEMENT DERIVATION: DDARKBRT is computed using the following algorithm:

```
DDARKBRT = DURATION;
IF DURATION
  THEN DDARKBRT = DDARKBCT/DURATION;
```

RANGE OF VALUE: DDARKBRT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: DDARKBRT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

```
+-----------------+
| D D A S Y S C T |  System Transfers
+-----------------+
DDASYSCT contains the number of physical disk reads and
physical disk writes performed for file system management.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  DDADMMCT - Memory Manager Transfers

ELEMENT DERIVATION:  DDASYSCT is set to the value of the
MeasureWare Disk metric BYDSK_SYSTEM_IO.

RANGE OF VALUE:  DDASYSCT takes any numeric value ranging
from 0 to a maximum that is limited by practical bounds.

SUMMARIZATION PROCESS:  DDASYSCT contains the sum of all
values of DDASYSCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| D D A W K B C T |  KBytes Written
+-----------------+
DDAWKBCT contains the number of KBytes transferred to the
disk device during the measurement interval.  This counts all
types of disk writes, including file system, virtual memory,
and raw IO.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  DDARKBCT - KBytes Read
           DDAWKBRT - Disk Writer Transfer Rate in KBytes

ELEMENT DERIVATION:  DDAWKBCT is set to the value of the
MeasureWare Application metric bydsk_PHYS_WRITE_BYTE.

RANGE OF VALUE:  DDAWKBCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAWKBCT contains the sum of all
values of DDAWKBCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| D D A W K B R T |  Disk Write Transfer Rate in KBytes
+-----------------+
DDAWKBRT contains the average number of KBytes per second
transferred to the disk device during the measurement
interval.  This count includes all types of writes, including
file system, virtual memory, and raw IO.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DDAWKBCT - KBytes Written
           DURATION - Recording Interval Time

ELEMENT DERIVATION:  DDAWKBRT is computed using the following
algorithm:

  DDAWKBRT = DURATION;
  IF DURATION
    THEN DDAWKBRT = DDAWKBCT/DURATION;

RANGE OF VALUE:  DDAWKBRT ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  DDAWKBRT contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None
```

B.8 - The UNIX System and Collector Profile (AUMMCF) File

The UNIX System and Collector Profile contains global system measurements useful for problem diagnosis, change tracking, and reporting. It includes information about the UNIX system on which data was collected, and is derived from MeasureWare Configuration metrics.

The metrics supported vary by operating system.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	YEAR	MONTH	DAY
	HOUR	ENDTS			
DAYS	N/A				
WEEKS	N/A				
MONTHS	SYSID	AUMOPSYS	YEAR	MONTH	ZONE
YEARS	N/A				
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.

NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

- X - DETAIL
- D - DAYS
- W - WEEKS
- M - MONTHS
- Y - YEARS
- T - TABLES AREA
- . - File is not supported

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
Sequence/Summary Data Elements		
X..M..E	AUMOPSYS	- Operating System Name
X.....E	DAY	- Day of Month
X.....E	HOUR	- Hour of Day
X..M..E	MONTH	- Month of Year
X..M..E	SYSID	- System Identifier
X.....E	WEEK	- Week of Year
X..M..E	YEAR	- Year of Century
X..M..E	ZONE	- Time Zone
Common Data Elements		

```
 X..M..E  AUMCPMDL - System Model Identification
 X..M..E  AUMLGVER - Measurement Interface Version Number
 X..M..E  AUMOPSLC - Operating System License
 X..M..E  AUMOPSVR - Operating System Release Number
 X..M..E  AUMSCOPE - Scope Collector Name and Version Number
 X..M..E  AUMSYSID - Name of Host System
 X.....E  DAYNAME  - Name of Day of Week
 X..M..E  ENDTS    - End Time Stamp
 X.....E  ORGSYSID - Originating System Identification

Retained Data Elements

 X..M..E  MCFCACHE - Size of Buffer Cache Configured
 X..M..E  MCFCPUMN - Threshold, Percent of CPU
 X..M..E  MCFINODE - Entries in Inode Table on System
 X..M..E  MCFIORMN - Threshold, Physical Disk I/O Rate
 X..M..E  MCFKRNSZ - Kernel Word Size
 X..M..E  MCFLCKAV - File/Record Locks Available
 X..M..E  MCFLGTPS - Data Types Logged by Collector
 X..M..E  MCFLOSHR - Flag - Log Short-Running Processes
 X..M..E  MCFMCMOD - CPU Model
 X..M..E  MCFMEMSW - Swap Space Configured
 X..M..E  MCFMEMSZ - Main Memory Configured
 X..M..E  MCFMEMUS - Main Memory Available for Users
 X..M..E  MCFMSGQU - Message Queues Available
 X..M..E  MCFNCPUS - Number of Processors Configured
 X..M..E  MCFNDSKS - Number of Disk Drives Configured
 X..M..E  MCFNLANS - Number of LAN Interfaces Configured
 X..M..E  MCFNOKIL - Flag - No Terminating Processes
 X..M..E  MCFNONEW - Flag - No Newly Created Processes
 X..M..E  MCFNPROC - Value of NPROC Parameter
 X..M..E  MCFOPSYS - Operating System Name
 X..M..E  MCFPRCMM - Threshold, Virtual Memory
 X..M..E  MCFSEMID - Semaphore Identifiers Available
 X..M..E  MCFSHMSG - Table Memory Segments Available
 X..M..E  MCFTBLUK - Entries Used by Kernel
 X..M..E  MCFWCPMN - Threshold, Percent of CPU Wait
 X..M..E  MCFWDKMN - Threshold, Percent of Disk Wait
 X..M..E  MCFWIMMN - Threshold, Percent of Semaphore Wait
 X..M..E  MCFWMEMN - Threshold, Percent of Memory Wait

Minimum Data Elements

 X..M..E  MCFMNCHC - Minimum Size of Buffer Cache Configured
 X..M..E  MCFMNCPU - Minimum Number of Processors Configured
 X..M..E  MCFMNDSK - Minimum Number of Disk Drives Configured
 X..M..E  MCFMNLAN - Minimum Num of LAN Interfaces Configured
 X..M..E  MCFMNMEM - Minimum Main Memory Configured
 X..M..E  MCFMNMUS - Minimum Main Memory Available for Users
 X..M..E  MCFMNSWS - Minimum Swap Space Configured

Maximum Data Elements

 X..M..E  MCFMXCHC - Maximum Size of Buffer Cache Configured
 X..M..E  MCFMXCPU - Maximum Number of Processors Configured
 X..M..E  MCFMXDSK - Maximum Number of Disk Drives Configured
 X..M..E  MCFMXLAN - Maximum Num of LAN Interfaces Configured
 X..M..E  MCFMXMEM - Maximum Main Memory Configured
 X..M..E  MCFMXMUS - Maximum Main Memory Available for Users
 X..M..E  MCFMXSWS - Maximum Swap Space Configured
             File

+-----------------+
| M C F C A C H E |  Size of Cache Buffer Configured
+-----------------+
MCFCACHE contains the number of kilobytes configured for the
cache buffer.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  MCFMEMSW - Swap Space Configured
           MCFMEMSZ - Main Memory Configured

ELEMENT DERIVATION:  MCFCACHE is set to the value of the
MeasureWare Configuration metric TBL_BUFFER_CACHE_AVAIL.

RANGE OF VALUE:  MCFCACHE takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  MCFCACHE contains the last value of
MCFCACHE processed in the level of summarization.

USAGE NOTES:  None

+-----------------+
| M C F C P U M N |  Threshold, Percent of CPU
+-----------------+
MCFCPUMN contains the amount of CPU that a process must use
to become interesting.

TYPE:  Retained - Numeric, measurement unit is real number.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  MCFCPUMN is set to the value of the
MeasureWare Configuration metric GBL_THRESHOLD_CPU.

RANGE OF VALUE:  MCFCPUMN takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.
```

SUMMARIZATION PROCESS: MCFCPUMN contains the last value of MCFCPUMN processed in the level of summarization.

USAGE NOTES: MCFCPUMN is expressed as a percentage of elapsed time.

M C F I N O D E — Entries in Inode Table on System

MCFINODE contains the number of entries for the inode tables on the system.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: MCFINODE is set to the value of the MeasureWare Configuration metric TBL_INODE_CACHE_AVAIL.

RANGE OF VALUE: MCFINODE ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFINODE contains the last value of MCFINODE processed in the level of summarization.

USAGE NOTES: None

M C F I O R M N — Threshold, Physical Disk I/O Rate

MCFIORMN contains the number of physical disk I/Os per second that a process must exceed to become interesting.

TYPE: Retained - Numeric, measurement unit is real number.

SEE ALSO: Not applicable

ELEMENT DERIVATION: MCFIORMN is set to the value of the MeasureWare Configuration metric GBL_THRESHOLD_DISK.

RANGE OF VALUE: MCFIORMN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFIORMN contains the last value of MCFIORMN processed in the level of summarization.

USAGE NOTES: None

M C F K R N S Z — Kernel Word Size

MCFKRNSZ contains the word size of the current kernel on the system.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: MCFKRNSZ is set to the value of the MeasureWare Configuration metric GBL_OSKERNELTYPE_INT.

RANGE OF VALUE: MCFKRNSZ ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFKRNSZ contains the last value of MCFKRNSZ processed in the level of summarization.

USAGE NOTES: None

M C F L C K A V — File/Record Locks Available

MCFLCKAV contains the number of record or file locks that may be allocated on the system.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: MCFLCKAV is set to the value of the MeasureWare Configuration metric TBL_FILE_LOCK_AVAIL.

RANGE OF VALUE: MCFLCKAV ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFLCKAV contains the last value of MCFLCKAV processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| M C F L G T P S |  Data Types Logged by Collector
+-----------------+
MCFLGTPS contains codes that signify what kind of data is
being logged by the collector.

The default code for the HP-UX platform is "G P", meaning
that global and interesting process data is being logged.

The default code for the AIX and SunOS platforms is "G PD",
meaning that global and interesting processes, as well as
disk device data, is being logged.

TYPE:  Retained - Character, length 16.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  MCFLGTPS is set to the value of the
MeasureWare Configuration metric GBL_LOGGING_TYPES.

RANGE OF VALUE:  MCFLGTPS can contain any combination of 1 to
8 alphanumeric characters.

SUMMARIZATION PROCESS:  MCFLGTPS contains the last value of
MCFLGTPS processed in the level of summarization.

USAGE NOTES:  To interpret the coded values, consult your
MeasureWare documentation for the release of MeasureWare that
your installation is running (CA MICS element AUMSCOPE).


+-----------------+
| M C F L O S H R |  Flag - Log Short-Running Processes
+-----------------+
MCFLOSHR contains a flag specifying that short-running
processes (a process that runs for less than one second)
should be logged as interesting when they terminate.

TYPE:  Retained - Character, length 8.

SEE ALSO:  MCFNOKIL - Flag, No Terminating Processes
           MCFNONEW - Flag, No Newly Created Processes

ELEMENT DERIVATION:  MCFLOSHR is set to the value of the
MeasureWare Configuration metric GBL_THRESHOLD_SHORTLIVED.

RANGE OF VALUE:  MCFLOSHR may contain any combination of 1 to
8 alphanumeric characters.

SUMMARIZATION PROCESS:  MCFLOSHR contains the last value of
MCFLOSHR processed in the level of summarization.

USAGE NOTES:  MeasureWare by default does not log
short-running processes.


+-----------------+
| M C F M C M O D |  CPU Model
+-----------------+
MCFMCMOD represents the CPU model.

TYPE:  Retained - Character, length 20.

SEE ALSO:  AUMCPMDL - System Model Identification

ELEMENT DERIVATION:  MCFMCMOD is set to the value of the
MeasureWare Configuration metric GBL_MACHINE_MODEL.

RANGE OF VALUE:  MCFMCMOD takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  MCFMCMOD contains the last value of
MCFMCMOD processed in the level of summarization.

USAGE NOTES:  For the HP-UX platform, this data element
returns more information on processors such as the T500
series. On a Solaris platform, more information is returned
from processors such as the SUNW,Ultra-1.


+-----------------+
| M C F M E M S W |  Swap Space Configured
+-----------------+
MCFMEMSW contains the amount of potential swap space (in
kilobytes) and is updated every 30 seconds or at the
customized interval time, whichever is greater.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  MCFMEMSZ - Main Memory Configured
           MCFMEMUS - Main Memory Available for Users

ELEMENT DERIVATION:  MCFMEMSW is set to the value of the
MeasureWare Configuration metric GBL_SWAP_SPACE_AVAIL_KB.

RANGE OF VALUE:  MCFMEMSW takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  MCFMEMSW contains the last value of
MCFMEMSW processed in the level of summarization.
```

USAGE NOTES: None

| M C F M E M S Z | Main Memory Configured

MCFMEMSZ contains the amount of main memory (in kilobytes) configured.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: MCFMEMSW - Swap Space Configured
MCFMEMUS - Main Memory Available for Users

ELEMENT DERIVATION: MCFMEMSZ is set to the value of the MeasureWare Configuration metric GBL_MEM_PHYS.

RANGE OF VALUE: MCFMEMSZ takes any numeric value from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMEMSZ contains the last value of MCFMEMSZ processed in the level of summarization.

USAGE NOTES: None

| M C F M E M U S | Main Memory Available for Users

MCFMEMUS contains the amount of main memory (in kilobytes) available for user processes.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: MCFMEMSW - Swap Space Configured
MCFMEMSZ - Main Memory Configured

ELEMENT DERIVATION: MCFMEMUS is set to the value of the MeasureWare Configuration metric GBL_MEM_AVAIL.

RANGE OF VALUE: MCFMEMUS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMEMUS contains the last value of MCFMEMUS processed in the level of summarization.

USAGE NOTES: None

| M C F M N C H C | Minimum Size of Cache Buffer Configured

MCFMNCHC contains the minimum size of the cache buffer configured.

TYPE: Minimum - Numeric, measurement unit is integer.

SEE ALSO: MCFCACHE - Size of Cache Buffer Configured

ELEMENT DERIVATION: MCFMNCHC is set to MCFCACHE.

RANGE OF VALUE: MCFMNCHC takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMNCHC contains the minimum value of MCFMNCHC processed for the level of summarization.

USAGE NOTES: None

| M C F M N C P U | Minimum Number of Processors Configured

MCFMNCPU contains the minimum number of processors configured.

TYPE: Minimum - Numeric, measurement unit is integer.

SEE ALSO: MCFNCPUS - Number of Processors Configured

ELEMENT DERIVATION: MCFMNCPU is set to MCFNCPUS.

RANGE OF VALUE: MCFMNCPU takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMNCPU contains the minimum value of MCFMNCPU processed for the level of summarization.

USAGE NOTES: None

| M C F M N D S K | Minimum Number of Disk Drives Configured

MCFMNDSK contains the minimum number of disk drives configured.

TYPE: Minimum - Numeric, measurement unit is integer.

SEE ALSO: MCFNDSKS - Number of Disk Drives Configured

ELEMENT DERIVATION: MCFMNDSK is set to MCFNDSKS.

RANGE OF VALUE: MCFMNDSK takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMNDSK contains the minimum value of MCFMNDSK processed for the level of summarization.

USAGE NOTES: None

M C F M N L A N — Minimum Num of LAN Interfaces Configured

MCFMNLAN contains the minimum number of LAN interfaces configured.

TYPE: Minimum - Numeric, measurement unit is integer.

SEE ALSO: MCFNLANS - Number of LAN Interfaces Configured

ELEMENT DERIVATION: MCFMNLAN is set to the value of MCFNLANS.

RANGE OF VALUE: MCFMNLAN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMNLAN contains the minimum value of MCFMNLAN processed for the level of summarization.

USAGE NOTES: None

M C F M N M E M — Minimum Main Memory Configured

MCFMNMEM contains the minimum amount of main memory configured.

TYPE: Minimum - Numeric, measurement unit is integer.

SEE ALSO: MCFMEMSZ - Main Memory Configured

ELEMENT DERIVATION: MCFMNMEM is set to MCFMEMSZ.

RANGE OF VALUE: MCFMNMEM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMNMEM contains the minimum value of MCFMNMEM processed for the level of summarization.

USAGE NOTES: None

M C F M N M U S — Minimum Main Memory Available for Users

MCFMNMUS contains the minimum main memory available for users.

TYPE: Minimum - Numeric, measurement unit is integer.

SEE ALSO: MCFMEMUS - Main Memory Available for Users

ELEMENT DERIVATION: MCFMNMUS is set to MCFMEMUS.

RANGE OF VALUE: MCFMNMUS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMNMUS contains the minimum value of MCFMNMUS processed for the level of summarization.

USAGE NOTES: None

M C F M N S W S — Minimum Swap Space Configured

MCFMNSWS contains the minimum swap space configured.

TYPE: Minimum - Numeric, measurement unit is integer.

SEE ALSO: MCFMEMSW - Swap Space Configured

ELEMENT DERIVATION: MCFMNSWS is set to MCFMEMSW.

RANGE OF VALUE: MCFMNSWS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMNSWS contains the minimum value of MCFMNSWS processed for the level of summarization.

USAGE NOTES: None

```
+-----------------+
| M C F M S G Q U |  Message Queues Available
+-----------------+
MCFMSGQU contains the number of the message queues that can
be allocated on a system during the measurement interval.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  MCFMSGQU is set to the value of the
MeasureWare Configuration metric TBL_MSG_TABLE_AVAIL.

RANGE OF VALUE:  MCFMSGQU ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  MCFMSGQU contains the last value of
MCFMSGQU processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| M C F M X C H C |  Maximum Size of Cache Buffer Configured
+-----------------+
MCFMXCHC contains the maximum size of cache buffer
configured.

TYPE:  Maximum - Numeric, measurement unit is integer.

SEE ALSO:  MCFCACHE - Size of Cache Buffer Configured

ELEMENT DERIVATION:  MCFMXCHC is set to MCFCACHE.

RANGE OF VALUE:  MCFMXCHC takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  MCFMXCHC contains the maximum value
of MCFMXCHC processed for the level of summarization.

USAGE NOTES:  None


+-----------------+
| M C F M X C P U |  Maximum Number of Processors Configured
+-----------------+
MCFMXCPU contains the maximum number of processors
configured.

TYPE:  Maximum - Numeric, measurement unit is integer.

SEE ALSO:  MCFNCPUS - Number of Processors Configured

ELEMENT DERIVATION:  MCFMXCPU is set to MCFNCPUS.

RANGE OF VALUE:  MCFMXCPU takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  MCFMXCPU contains the maximum value
of MCFMXCPU processed for the level of summarization.

USAGE NOTES:  None


+-----------------+
| M C F M X D S K |  Maximum Number of Disk Drives Configured
+-----------------+
MCFMXDSK contains the maximum number of disk drives
configured.

TYPE:  Maximum - Numeric, measurement unit is integer.

SEE ALSO:  MCFNDSKS - Number of Disk Drives Configured

ELEMENT DERIVATION:  MCFMXDSK is set to MCFNDSKS.

RANGE OF VALUE:  MCFMXDSK takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  MCFMXDSK contains the maximum value
of MCFMXDSK processed for the level of summarization.

USAGE NOTES:  None


+-----------------+
| M C F M X L A N |  Maximum Num of LAN Interfaces Configured
+-----------------+
MCFMXLAN contains the maximum number of LAN interfaces
configured.

TYPE:  Maximum - Numeric, measurement unit is integer.

SEE ALSO:  MCFNLANS - Number of LAN Interfaces Configured

ELEMENT DERIVATION:  MCFMXLAN is set to MCFNLANS.

RANGE OF VALUE:  MCFMXLAN takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.
```

SUMMARIZATION PROCESS: MCFMXLAN contains the maximum value of MCFMXLAN processed for the level of summarization.

USAGE NOTES: None

| M C F M X M E M | Maximum Main Memory Configured

MCFMXMEM contains the maximum main memory configured.

TYPE: Maximum - Numeric, measurement unit is integer.

SEE ALSO: MCFMEMSZ - Main Memory Configured

ELEMENT DERIVATION: MCFMXMEM is set to MCFMEMSZ.

RANGE OF VALUE: MCFMXMEM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMXMEM contains the maximum value of MCFMXMEM processed for the level of summarization.

USAGE NOTES: None

| M C F M X M U S | Maximum Main Memory Available for Users

MCFMXMUS contains the maximum main memory available for users.

TYPE: Maximum - Numeric, measurement unit is integer.

SEE ALSO: MCFMEMUS - Main Memory Available for Users

ELEMENT DERIVATION: MCFMXMUS is set to MCFMEMUS.

RANGE OF VALUE: MCFMXMUS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMXMUS contains the maximum value of MCFMXMUS processed for the level of summarization.

USAGE NOTES: None

| M C F M X S W S | Maximum Swap Space Configured

MCFMXSWS contains the maximum swap space configured.

TYPE: Maximum - Numeric, measurement unit is integer.

SEE ALSO: MCFMEMSW - Swap Space Configured

ELEMENT DERIVATION: MCFMXSWS is set to MCFMEMSW.

RANGE OF VALUE: MCFMXSWS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFMXSWS contains the maximum value of MCFMXSWS processed for the level of summarization.

USAGE NOTES: None

| M C F N C P U S | Number of Processors Configured

MCFNCPUS contains the number of active processors configured.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: MCFNDSKS - Number of Disk Drives Configured
MCFNLANS - Number of LAN Interfaces Configured

ELEMENT DERIVATION: MCFNCPUS is set to the value of the MeasureWare Configuration metric GBL_NUM_CPU for HP-UX, or GBL_ACTIVE_CPU for Solaris.

RANGE OF VALUE: MCFNCPUS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFNCPUS contains the last value of MCFNCPUS processed in the level of summarization.

USAGE NOTES: None

| M C F N D S K S | Number of Disk Drives Configured

MCFNDSKS contains the number of active disk drives configured.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: MCFNCPUS - Number of Processors Configured
MCFNLANS - Number of LAN Interfaces Configured

ELEMENT DERIVATION: MCFNDSKS is set to the value of the MeasureWare Configuration metric GBL_NUM_DISK.

RANGE OF VALUE: MCFNDSKS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFNDSKS contains the last value of MCFNDSKS processed in the level of summarization.

USAGE NOTES: None

M C F N L A N S — Number of LAN Interfaces Configured

MCFNLANS contains the number of LAN interfaces configured.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: MCFNCPUS - Number of Processors Configured
MCFNDSKS - Number of Disk Drives Configured

ELEMENT DERIVATION: MCFNLANS is set to the value of the MeasureWare Configuration metric GBL_NUM_NETWORK.

RANGE OF VALUE: MCFNLANS takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFNLANS contains the last value of MCFNLANS processed in the level of summarization.

USAGE NOTES: None

M C F N O K I L — Flag - No Terminating Processes

MCFNOKIL contains a flag specifying that terminating processes are not interesting unless they exceed at least one other threshold.

TYPE: Retained - Character, length 12.

SEE ALSO: MCFNONEW - Flag, No Newly Created Processes
MCFLOSHR - Flag, Log Short-Running Processes

ELEMENT DERIVATION: MCFNOKIL is set to the value of the MeasureWare Configuration metric GBL_THRESHOLD_NOKILLED.

RANGE OF VALUE: MCFNOKIL can contain any combination of 1 to 8 alphanumeric characters.

SUMMARIZATION PROCESS: MCFNOKIL contains the last value of MCFNOKIL processed in the level of summarization.

USAGE NOTES: MeasureWare normally logs processes that terminate. If your MeasureWare system has this flag set, the amount of data in your system's log files is minimized; therefore, the amount of data transferred to MVS for CA MICS is minimized.

M C F N O N E W — Flag - No Newly Created Processes

MCFNONEW contains a flag specifying that newly created processes are not interesting.

TYPE: Retained - Character, length 8.

SEE ALSO: MCFNOKIL - Flag, No Terminating Processes
MCFLOSHR - Flag, Log Short-Running Processes

ELEMENT DERIVATION: MCFNONEW is set to the value of the MeasureWare Configuration metric GBL_THRESHOLD_NONEW.

RANGE OF VALUE: MCFNONEW can contain any combination of 1 to 8 alphanumeric characters.

SUMMARIZATION PROCESS: MCFNONEW contains the last value of MCFNONEW processed in the level of summarization.

USAGE NOTES: MeasureWare normally logs new processes. If your MeasureWare system specified this flag, the amount of data in your system's log files is minimized and therefore the amount of data transferred to MVS for CA MICS is minimized.

M C F N P R O C — Value of NPROC Parameter

MCFNPROC contains the size of the process table.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: MCFNPROC is set to the value of the MeasureWare Configuration metric TBL_PROC_TABLE_AVAIL.

RANGE OF VALUE: MCFNPROC takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFNPROC contains the last value of MCFNPROC processed in the level of summarization.

USAGE NOTES: On the HP-UX platform, MCFNPROC is set by the NPROC value during system generation.

M C F O P S Y S — Operating System Name

MCFOPSYS contains the operating system name (HP-UX, for example).

TYPE: Common, retained - Character, length 12.

SEE ALSO: AUMOPSYS - Operating System Name

ELEMENT DERIVATION: MCFOPSYS is set to the value of the MeasureWare Configuration metric GBL_OSNAME.

RANGE OF VALUE: MCFOPSYS can contain any combination of 1 to 12 alphanumeric characters.

SUMMARIZATION PROCESS: MCFOPSYS contains the last value of MCFOPSYS processed in the level of summarization.

USAGE NOTES: None

M C F P R C M M — Threshold, Virtual Memory

MCFPRCMM contains the amount of virtual memory (in megabytes) that a process must use to become interesting during an interval.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: MCFPRCMM is set to the value of the MeasureWare Configuration metric GBL_THRESHOLD_PROCMEM.

RANGE OF VALUE: MCFPRCMM ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFPRCMM contains the last value of MCFPRCMM processed in the level of summarization.

USAGE NOTES: In order for a process to be considered "interesting" and logged, it must exceed at least one threshold value in any given interval.

M C F S E M I D — Semaphore Identifiers Available

MCFSEMID contains the number of semaphore identifiers that can be allocated on the system.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: MCFSEMID is set to the value of the MeasureWare Configuration metric TBL_SEM_TABLE_AVAIL.

RANGE OF VALUE: MCFSEMID ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFSEMID contains the last value of MCFSEMID processed in the level of summarization.

USAGE NOTES: None

M C F S H M S G — Table Memory Segments Available

MCFSHMSG contains the number of shared memory segments that can be allocated on a system during the measurement interval.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: MCFSHMSG is set to the value of the MeasureWare Configuration metric TBL_SHMEM_TABLE_AVAIL.

RANGE OF VALUE: MCFSHMSG ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFSHMSG contains the last value of MCFSHMSG processed in the level of summarization.

USAGE NOTES: None

| M C F T B L U K | Entries Used by Kernel

MCFTBLUK is the number of file table entries used by the kernel to manage open file descriptors.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: None

ELEMENT DERIVATION: MCFTBLUK is set to the value of the MeasureWare Configuration metric TBL_FILE_TABLE_AVAIL.

RANGE OF VALUE: MCFTBLUK takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFTBLUK contains the last value of MCFTBLUK processed in the level of summarization.

USAGE NOTES: None

| M C F W C P M N | Threshold, Percent of CPU Wait

MCFWCPMN contains the percentage of time a process waits to become interesting during a measurement interval.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: MCFWDKMN - Threshold, Percent of Disk Wait
MCFWMEMN - Threshold, Percent of Memory Wait
MCFWIMMN - Threshold, Percent of Semaphore Wait

ELEMENT DERIVATION: MCFWCPMN is set to the value of the MeasureWare Configuration metric GBL_THRESHOLD_WAIT_CPU.

RANGE OF VALUE: MCFWCPMN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFWCPMN contains the last value of MCFWCPMN processed in the level of summarization.

USAGE NOTES: None

| M C F W D K M N | Threshold, Percent of Disk Wait

MCFWDKMN contains the percentage of time a process waits for a disk subsystem to become interesting during a measurement interval.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: MCFWCPMN - Threshold, Percent of CPU Wait
MCFWMEMN - Threshold, Percent of Memory Wait
MCFWIMMN - Threshold, Percent of Semaphore Wait

ELEMENT DERIVATION: MCFWDKMN is set to the value of the MeasureWare Configuration metric GBL_THRESHOLD_WAIT_DISK.

RANGE OF VALUE: MCFWDKMN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFWDKMN contains the last value of MCFWDKMN processed in the level of summarization.

USAGE NOTES: None

| M C F W I M M N | Threshold, Percent of Semaphore Wait

MCFWIMMN contains the percentage of time a process spent waiting for a semaphore to become interesting during a measurement interval.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: MCFWDKMN - Threshold, Percent of Disk Wait
MCFWMEMN - Threshold, Percent of Memory Wait
MCFWCPMN - Threshold, Percent of CPU Wait

ELEMENT DERIVATION: MCFWIMMN is set to the value of the MeasureWare Configuration metric GBL_THRESHOLD_WAIT_IMPEDE.

RANGE OF VALUE: MCFWIMMN takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: MCFWIMMN contains the last value of MCFWIMMN processed in the level of summarization.

USAGE NOTES: None

```
+-----------------+
| M C F W M E M N |  Threshold, Percent of Memory Wait
+-----------------+
MCFWMEMN contains the percentage of time a process spent
waiting to access main memory to become interesting during a
measurement interval.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  MCFWDKMN - Threshold, Percent of Disk Wait
           MCFWIMMN - Threshold, Percent of Semaphore Wait
           MCFWCPMN - Threshold, Percent of CPU Wait

ELEMENT DERIVATION:  MCFWMEMN is set to the value of the
MeasureWare Configuration metric GBL_THRESHOLD_WAIT_MEMORY.

RANGE OF VALUE:  MCFWMEMN takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  MCFWMEMN contains the last value of
MCFWMEMN processed in the level of summarization.

USAGE NOTES:  None
```

B.9 - The UNIX Logical Volume Activity (AUMVLM) File

The UNIX Logical Volume Activity file contains interval-based measurements useful for capacity analyses. It includes information about I/O and space for individual logical volumes. It is derived from MeasureWare Volume metrics, which are available from HP-UX systems.

The metrics supported vary by operating system.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity				
DETAIL	SYSID	AUMOPSYS	VLMNAME	YEAR	MONTH
	DAY	HOUR	ENDTS		
DAYS	SYSID	AUMOPSYS	VLMNAME	YEAR	MONTH
	DAY	HOUR			
WEEKS	SYSID	AUMOPSYS	VLMNAME	YEAR	WEEK
	ZONE	HOUR			
MONTHS	SYSID	AUMOPSYS	VLMNAME	YEAR	MONTH
	ZONE				
YEARS	SYSID	AUMOPSYS	VLMNAME	YEAR	ZONE
TABLES	N/A				

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.
NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

```
X - DETAIL
D - DAYS
W - WEEKS
M - MONTHS
Y - YEARS
T - TABLES AREA
. - File is not supported
```

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
		Sequence/Summary Data Elements
XDWMY.E	AUMOPSYS	- Operating System Name
XD....E	DAY	- Day of Month
XDW...E	HOUR	- Hour of Day
XD.M..E	MONTH	- Month of Year
XDWMY.E	SYSID	- System Identifier
XDWMY.E	VLMNAME	- Logical Volume Group Name
XDW...E	WEEK	- Week of Year
XDWMY.E	YEAR	- Year of Century
XDWMY.E	ZONE	- Time Zone
		Common Data Elements
XD....E	DAYNAME	- Name of Day of Week
XDWMY.E	DURATION	- Recording Interval Time
XDWMY.E	ENDTS	- End Time Stamp
XDWMY.E	INTERVLS	- Number of Recording Intervals
XDWMY.E	STARTTS	- Start Time Stamp
		Retained Data Elements
XDWMY.E	VLMDIR	- Logical Volume Directory
XDWMY.E	VLMPTHNM	- Mount Point Path Name
		Accumulated Data Elements
XDWMY.E	VLMLRDCT	- Logical Disk Reads
XDWMY.E	VLMLWTCT	- Logical Disk Writes
XDWMY.E	VLMSPUSM	- Sum of Avg Disk Space Used
		Minimum Data Elements
XDWMY.E	VLMMNSPU	- Minimum Pct Disk Space Used
		Maximum Data Elements
XDWMY.E	VLMMXSPU	- Maximum Pct Disk Space Used
		Derived Data Elements
XDWMY.E	VLMPCSPU	- Percent Disk Space Used
XDWMY.E	VLMPSLRD	- Logical Disk Read Rate
XDWMY.E	VLMPSLWT	- Logical Disk Write Rate

V L M D I R — Logical Volume Directory

VLMDIR contains the path name of the logical volume or volume group during the measurement interval.

TYPE: Retained - Character, length 40.

SEE ALSO: VLMNAME - Logical Volume Group Name

ELEMENT DERIVATION: VLMDIR is set to the value of MeasureWare LVolume metric LV_DIRNAME.

RANGE OF VALUE: VLMDIR can contain any combination of 1 to 40 alphanumeric characters.

SUMMARIZATION PROCESS: VLMDIR is not summarized. VLMDIR is used to determine a unique key for records within a file

based on the timespan, so that other elements in the file can be summarized.

USAGE NOTES: None

| V L M L R D C T | Logical Disk Reads

VLMLRDCT contains the number of logical disk reads from this logical volume during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: VLMLWTCT - Logical Disk Writes
VLMSPUSM - Sum of Avg Disk Space Used

ELEMENT DERIVATION: VLMLRDCT is set to the value of the MeasureWare LVolume metric LV_LOGL_READ.

RANGE OF VALUE: VLMLRDCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: VLMLRDCT contains the sum of all values of VLMLRDCT processed in the level of summarization.

USAGE NOTES: None

| V L M L W T C T | Logical Disk Writes

VLMLWTCT contains the number of logical disk writes to this logical volume during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: VLMLRDCT - Logical Disk Reads
VLMSPUSM - Sum of Avg Disk Space Used

ELEMENT DERIVATION: VLMLWTCT is set to the value of the MeasureWare LVolume metric LV_LOGL_WRITE.

RANGE OF VALUE: VLMLWTCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: VLMLWTCT contains the sum of all values of VLMLWTCT processed in the level of summarization.

USAGE NOTES: None

| V L M M N S P U | Minimum Pct Disk Space Used

VLMMNSPU contains the minimum percentage of disk space used on the logical volume during the measurement interval.

TYPE: Minimum - Numeric, measurement unit is real number.

SEE ALSO: VLMPCSPU - Percent Disk Space Used

ELEMENT DERIVATION: At the DETAIL time-span, VLMMNSPU is set to the minimum value of the MeasureWare LVolume metric LV_SPACE_UTIL.

RANGE OF VALUE: The value of VLMMNSPU ranges from a minimum of 0 to a maximum of 100.

SUMMARIZATION PROCESS: VLMMNSPU contains the minimum value of VLMMNSPU processed for the level of summarization.

USAGE NOTES: None

| V L M M X S P U | Maximum Pct Disk Space Used

VLMMXSPU contains the maximum percentage of disk space used on the logical volume during the measurement interval.

TYPE: Maximum - Numeric, measurement unit is real number.

SEE ALSO: VLMMNSPU - Minimum Pct Disk Space Used
VLMPCSPU - Average Percent Disk Space Used

ELEMENT DERIVATION: VLMMXSPU is set to the maximum value of VLMPCSPU processed for the level of summarization.

RANGE OF VALUE: The value of VLMMXSPU ranges from a minimum of 0 to a maximum of 100.

SUMMARIZATION PROCESS: VLMMXSPU contains the maximum value of VLMMXSPU processed for the level of summarization.

USAGE NOTES: None

```
+-----------------+
| V L M N A M E   |  Logical Volume Group Name
+-----------------+
VLMNAME contains the name of the volume group associated
with the logical volume during the measurement interval.

TYPE:  Retained - Character, length 40.

SEE ALSO:  VLMPDIR  - Logical Volume Directory

ELEMENT DERIVATION:  VLMNAME is set to the value of the
MeasureWare LVolume metric LV_GROUP_NAME.

RANGE OF VALUE:  VLMNAME can contain any combination of 1 to
40 alphanumeric characters.

SUMMARIZATION PROCESS:  VLMNAME is not summarized.  VLMNAME
is used to determine a unique key for records within a file
based on the time-span, so other elements in the file can be
summarized.

USAGE NOTES:  None


+-----------------+
| V L M P C S P U |  Percent Disk Space Used
+-----------------+
VLMPCSPU contains the average percentage of disk space used
on this logical volume during the measurement interval.

TYPE:  Computed using specific EXPression statements -
  Numeric, measurement unit is real number.

SEE ALSO:  INTERVLS - Number of Recording Intervals
           VLMSPUSM - Sum of Avg Disk Space Used

ELEMENT DERIVATION:  VLMPCSPU is computed using the following
algorithm:

  VLMPCSPU = INTERVLS;
  IF INTERVLS
    THEN VLMPCSPU = VLMSPUSM/INTERVLS;

RANGE OF VALUE:  The value of VLMPCSPU ranges from a minimum
of 0 to a maximum of 100.

SUMMARIZATION PROCESS:  VLMPCSPU contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| V L M P S L R D |  Logical Disk Read Rate
+-----------------+
VLMPSLRD contains the rate at which this logical volume was
read during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           VLMLRDCT - Logical Disk Reads

ELEMENT DERIVATION:  VLMPSLRD is computed using the following
algorithm:

  VLMPSLRD = DURATION;
  IF DURATION
    THEN VLMPSLRD = VLMLRDCT/DURATION;

RANGE OF VALUE:  VLMPSLRD ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  VLMPSLRD contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| V L M P S L W T |  Logical Disk Write Rate
+-----------------+
VLMPSLWT contains the rate to which this logical volume was
written during the measurement interval.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is real number.

SEE ALSO:  DURATION - Recording Interval Time
           VLMLWTCT - Logical Disk Writes

ELEMENT DERIVATION:  VLMPSLWT is computed using the following
algorithm:

  VLMPSLWT = DURATION;
  IF DURATION
    THEN VLMPSLWT = VLMLWTCT/DURATION;
```

RANGE OF VALUE: VLMPSLWT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: VLMPSLWT contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

V L M P T H N M | Mount Point Path Name

VLMPTHNM contains the path name of the mount point of the file system during the measurement interval.

TYPE: Retained - Character, length 40.

SEE ALSO: Not applicable

ELEMENT DERIVATION: VLMPTHNM is set to the value of the MeasureWare LVolume metric FS_DIRNAME.

RANGE OF VALUE: VLMPTHNM can contain any combination of 1 to 40 alphanumeric characters.

SUMMARIZATION PROCESS: VLMPTHNM contains the last value of VLMPTHNM processed in the level of summarization.

USAGE NOTES: None

V L M S P U S M | Sum of Avg Disk Space Used

VLMSPUSM contains the sum of the average amount of disk space used by this logical volume, that was sampled during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is real number.

SEE ALSO: VLMLRDCT - Logical Disk Reads
VLMLWTCT - Logical Disk Writes

ELEMENT DERIVATION: In the DETAIL time span, VLMSPUSM is computed using the following calculation:

VLMSPUSM = LV_SPACE_UTIL*AUMSAMPL

where LV_SPACE_UTIL is from LVolume data and AUMSAMPL is from the CA MICS database.

RANGE OF VALUE: VLMSPUSM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: VLMSPUSM contains the sum of all values of VLMSPUSM processed in the level of summarization.

USAGE NOTES: None

B.10 - The UNIX Bus. Transaction SLO Activity (AUMTTA) File

The UNIX Business Transaction SLO Activity (AUMTTA) file contains measurements useful for managing the service level objectives your business application users require. It includes information about transaction response time and how transactions perform within established SLOs. This file is derived from MeasureWare Transaction metrics.

The table below identifies data elements by which the file is sequenced and summarized in each timespan. N/A indicates that the file is not supported in a timespan. At the DETAIL level, data is sequenced but not summarized.

NOTE: The timespans in which a file is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

Timespan	Level of Data Granularity
DETAIL	N/A
DAYS	N/A
WEEKS	N/A
MONTHS	N/A
YEARS	N/A
TABLES	N/A

Generation Date: Thu, Jul 18, 2019

NOTE: This file was generated with ESSENTIAL=ALL option in effect. All data elements defined in the file are generated.
NOTE: This file was generated with DERIVED=DEFault option in effect. Whether data elements are kept on the file on auxiliary storage or not is controlled by the complex definition of the DERIVED option.

The table below identifies data elements contained in this file. The entries for each data element are:

TIMESPAN: Defines the timespans in which the data element is supported. The timespans are indicated by the letters "XDWMYT" as follows:

- X - DETAIL
- D - DAYS
- W - WEEKS
- M - MONTHS
- Y - YEARS
- T - TABLES AREA
- . - File is not supported

The timespan field also indicates Essential Elements with the letter E, if applicable.

DATA ELEMENT: The data element name.

DATA ELEMENT DESCRIPTION: The data element's long name.

The timespans in which a data element is supported are defined by each installation when CA MICS is installed. Therefore, this table has been generated as part of the installation process to accurately reflect the CA MICS system at your installation.

GENERATION DATE: Thu, Jul 18, 2019

Note: Essential data elements are identified by an "E" under the Timespan asterisk (*) column.

Time-Span *	Data Element	Data Element Description (LABEL)
Sequence/Summary Data Elements		
.....E	AUMOPSYS	- Operating System Name
.....E	DAY	- Day of Month
.....E	HOUR	- Hour of Day
.....E	MONTH	- Month of Year
.....E	SYSID	- System Identifier
.....E	TTANAME	- Transaction Name
.....E	WEEK	- Week of Year
.....E	YEAR	- Year of Century
.....E	ZONE	- Time Zone

```
Common Data Elements

 ......E  DAYNAME  - Name of Day of Week
 ......E  DURATION - Recording Interval Time
 ......E  ENDTS    - End Time Stamp
 ......E  INTERVLS - Number of Recording Intervals
 ......E  STARTTS  - Start Time Stamp

Retained Data Elements

 ......E  TTAADFMT - Correlator Address Format
 ......E  TTAARMID - ARM Transaction ID
 ......E  TTAARMIF - ARM Transaction Information
 ......E  TTAARMNM - ARM Application Name
 ......E  TTAARMUN - ARM Transaction User Name
 ......E  TTABINCT - Number of Transaction Bins Defined
 ......E  TTACORAD - Correlator Address
 ......E  TTASLOMX - Transaction Service Level Objective
 ......E  TTATFAIL - Number of Failed Transactions
 ......E  TTATRNID - Transaction Class ID
 ......E  TTAUPR01 - Upper Range Response Time for Bin 1
 ......E  TTAUPR02 - Upper Range Response Time for Bin 2
 ......E  TTAUPR03 - Upper Range Response Time for Bin 3
 ......E  TTAUPR04 - Upper Range Response Time for Bin 4
 ......E  TTAUPR05 - Upper Range Response Time for Bin 5
 ......E  TTAUPR06 - Upper Range Response Time for Bin 6
 ......E  TTAUPR07 - Upper Range Response Time for Bin 7
 ......E  TTAUPR08 - Upper Range Response Time for Bin 8
 ......E  TTAUPR09 - Upper Range Response Time for Bin 9
 ......E  TTAUPR10 - Upper Range Response Time for Bin 10

Accumulated Data Elements

 ......E  TTAABORT - Aborted Transaction Count
 ......E  TTAABTTM - Cumulative Aborted Transaction Time
 ......E  TTACPUTM - Total CPU Consumed by Transaction
 ......E  TTADSKIO - Physical Disk I/Os per Transaction
 ......E  TTALGLIO - Logical I/Os per Transaction
 ......E  TTATRANS - Completed Transaction Count
 ......E  TTATRNTM - Cumulative Transaction Response Time
 ......E  TTATRN01 - Completed Transactions in Bin 1
 ......E  TTATRN02 - Completed Transactions in Bin 2
 ......E  TTATRN03 - Completed Transactions in Bin 3
 ......E  TTATRN04 - Completed Transactions in Bin 4
 ......E  TTATRN05 - Completed Transactions in Bin 5
 ......E  TTATRN06 - Completed Transactions in Bin 6
 ......E  TTATRN07 - Completed Transactions in Bin 7
 ......E  TTATRN08 - Completed Transactions in Bin 8
 ......E  TTATRN09 - Completed Transactions in Bin 9
 ......E  TTATRN10 - Completed Transactions in Bin 10
 ......E  TTAVSLCT - Violated SLO Transaction Count
```

```
Derived Data Elements

 ......E  TTAAVATM - Average Time Per Aborted Trans
 ......E  TTAAVTTM - Average Transaction Response Time
 ......E  TTAPCB01 - Percent Transactions in Bin 1
 ......E  TTAPCB02 - Percent Transactions in Bin 2
 ......E  TTAPCB03 - Percent Transactions in Bin 3
 ......E  TTAPCB04 - Percent Transactions in Bin 4
 ......E  TTAPCB05 - Percent Transactions in Bin 5
 ......E  TTAPCB06 - Percent Transactions in Bin 6
 ......E  TTAPCB07 - Percent Transactions in Bin 7
 ......E  TTAPCB08 - Percent Transactions in Bin 8
 ......E  TTAPCB09 - Percent Transactions in Bin 9
 ......E  TTAPCB10 - Percent Transactions in Bin 10
 ......E  TTAPCVSL - Percent Transactions Violated SLO
              File
```

```
+-----------------+
| T T A A B O R T |  Aborted Transaction Count
+-----------------+
TTAABORT contains the number of transactions that aborted
during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  TTAABTTM - Cumulative Aborted Transaction Time
           TTAAVATM - Average Time Per Aborted Transaction

ELEMENT DERIVATION:  TTAABORT is set to the value of the
MeasureWare Transaction metric TT_ABORT.

RANGE OF VALUE:  TTAABORT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTAABORT contains the sum of all
values of TTAABORT processed in the level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| T T A A B T T M |  Cumulative Aborted Transaction Time
+-----------------+
TTAABTTM contains the cumulative wall clock time for aborted
transactions during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  TTAABORT - Aborted Transaction Count
           TTAAVATM - Average Time Per Aborted Transaction
```

ELEMENT DERIVATION: TTAABTTM is computed using MeasureWare metrics in the following calculation:

```
TT_ABORT_WALL_TIME_PER_TRAN*TT_ABORT
```

RANGE OF VALUE: TTAABTTM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: TTAABTTM contains the sum of all values of TTAABTTM processed in the level of summarization.

USAGE NOTES: None

T T A A D F M T | Correlator Address Format

TTAADFMT contains the correlator address format and shows the protocol family for the client network address.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: TTAADFMT is set to the value of the MeasureWare Transaction metric TT_CLIENT_ADDRESS_FORMAT.

RANGE OF VALUE: TTAADFMT ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: TTAADFMT contains the last value of TTAADFMT processed in the level of summarization.

USAGE NOTES: Refer to the HP ARM API Guide for a list and description of supported address formats.

T T A A R M I D | ARM Transaction ID

TTAARMID contains the Application Response Measurement (ARM) Transaction identifier for this transaction class.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: TTAARMNM - ARM Application Name
TTAARMUN - ARM Transaction User Name
TTAARMIF - ARM Transaction Information
TTANAME - Transaction Name

ELEMENT DERIVATION: TTAARMID is set to the value of the MeasureWare Transaction metric TT_TRAN_ID.

RANGE OF VALUE: TTAARMID ranges from a minimum of 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: TTAARMID contains the last value of TTAARMID processed in the level of summarization.

USAGE NOTES: None

T T A A R M I F | ARM Transaction Information

TTAARMIF contains information for this Application Response Measurement (ARM) transaction.

TYPE: Retained - Character, length 60.

SEE ALSO: TTAARMID - ARM Transaction ID
TTAARMNM - ARM Application Name
TTAARMUN - ARM Transaction User Name
TTANAME - Transaction Name

ELEMENT DERIVATION: TTAARMIF is set to the value of the MeasureWare Transaction metric TT_INFO.

RANGE OF VALUE: TTAARMIF can contain any combination of 1 to 60 alphanumeric characters.

SUMMARIZATION PROCESS: TTAARMIF contains the last value of TTAARMIF processed in the level of summarization.

USAGE NOTES: None

T T A A R M N M | ARM Application Name

TTAARMNM contains the name of the Application Response Measurement (ARM) application.

TYPE: Retained - Character, length 60.

SEE ALSO: TTAARMID - ARM Transaction ID
TTAARMNM - ARM Application Name
TTAARMUN - ARM Transaction User Name
TTAARMIF - ARM Transaction Information

```
            TTANAME  - Transaction Name

ELEMENT DERIVATION:  TTAARMNM is set to the value of the
MeasureWare Transaction metric TT_APP_NAME.

RANGE OF VALUE:  TTAARMNM can contain any combination of 1 to
60 alphanumeric characters.

SUMMARIZATION PROCESS:  TTAARMNM contains the last value of
TTAARMNM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A A R M U N |  ARM Transaction User Name
+-----------------+
TTAARMUN contains the user name of the Application Response
Measurement (ARM) transaction.

TYPE:  Retained - Character, length 60.

SEE ALSO:  TTAARMID - ARM Transaction ID
           TTAARMNM - ARM Application Name
           TTAARMIF - ARM Transaction Information
           TTANAME  - Transaction Name

ELEMENT DERIVATION:  TTAARMUN is set to the value of the
MeasureWare Transaction metric TT_UNAME.

RANGE OF VALUE:  TTAARMUN can contain any combination of 1 to
60 alphanumeric characters.

SUMMARIZATION PROCESS:  TTAARMUN contains the last value of
TTAARMUN processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A A V A T M |  Average Time Per Aborted Transaction
+-----------------+
TTAAVATM contains the average number of seconds for each
aborted transaction monitored by the Transaction Tracker.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is integer.

SEE ALSO:  TTAABTTM - Cumulative Aborted Transaction Time
           TTAABORT - Aborted Transaction Count


ELEMENT DERIVATION:  TTAAVATM is computed using the following
algorithm:

  TTAAVATM = TTAABORT;
  IF TTAABORT
    THEN TTAAVATM = TTAABTTM/TTAABORT;

RANGE OF VALUE:  TTAAVATM takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTAAVATM contains the result of the
computation described in ELEMENT DERIVATION. The value is
recomputed at each level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A A V T T M |  Average Transaction Response Time
+-----------------+
TTAAVTTM contains the average number of seconds for each
transaction monitored by the Transaction Tracker.

TYPE:  Computed using specific EXPression statements -
Numeric, measurement unit is integer.

SEE ALSO:  TTATRNTM - Cumulative Transaction Response Time
           TTATRANS - Completed Transaction Count

ELEMENT DERIVATION:  TTAAVTTM is computed using the following
algorithm:

  TTAAVTTM = TTATRANS;
  IF TTATRANS
    THEN TTAAVTTM = TTATRNTM/TTATRANS;

RANGE OF VALUE:  TTAAVTTM takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTAAVTTM contains the result of the
computation described in ELEMENT DERIVATION.  The value is
recomputed at each level of summarization.

USAGE NOTES:  None
```

```
+-----------------+
| T T A B I N C T |  Number of Transaction Bins Defined
+-----------------+
TTABINCT contains the number of Transaction Tracker
distribution ranges defined at your installation.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  TTABINCT is set to the value of the
MeasureWare Transaction metric TT_NUM_BINS.

RANGE OF VALUE:  TTABINCT takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTABINCT contains the last value of
TTABINCT processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A C O R A D |  Correlator Address
+-----------------+
TTACORAD contains the address where the child transaction
originated, the correlator address.

TYPE:  Retained - Character, length 20.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  TTACORAD is set to the value of the
MeasureWare Transaction metric TT_CLIENT_ADDRESS.

RANGE OF VALUE:  TTACORAD can contain any combination of 1 to
20 alphanumeric characters.

SUMMARIZATION PROCESS:  TTACORAD contains the last value of
TTACORAD processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A C P U T M |  Total CPU Consumed by Transaction
+-----------------+
TTACPUTM contains the total CPU time consumed by each
completed instance of the transactions during the measurement
interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  TTATRANS - Completed Transaction Count

ELEMENT DERIVATION:  TTACPUTM is computed using MeasureWare
Transaction metrics in the following calculation:

  TT_CPU_TOTAL_TIME_PER_TRAN*TTATRANS

RANGE OF VALUE:  TTACPUTM takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTACPUTM contains the sum of all
values of TTACPUTM processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A D S K I O |  Physical Disk I/Os per Transaction
+-----------------+
TTADSKIO contains the physical disk I/Os made by each
completed transaction during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  TTATRANS - Completed Transaction Count

ELEMENT DERIVATION:  TTADSKIO is computed using MeasureWare
Transaction metrics in the following calculation:

  TT_DISK_LOGL_IO_PER_TRAN*TT_ABORT

RANGE OF VALUE:  TTADSKIO takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTADSKIO contains the sum of all
values of TTADSKIO processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A L G L I O |  Logical I/Os per Transaction
+-----------------+
TTALGLIO contains the number of logical I/Os for each
completed instance of the transactions during the measurement
interval.
```

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: TTATRANS - Completed Transaction Count

ELEMENT DERIVATION: TTALGLIO is computed using MeasureWare Transaction metrics in the following calculation:

```
TT_DISK_PHYS_IO_PER_TRAN*TT_ABORT
```

RANGE OF VALUE: TTALGLIO takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: TTALGLIO contains the sum of all values of TTALGLIO processed in the level of summarization.

USAGE NOTES: None

T T A N A M E | Transaction Tracker Name

TTANAME contains the name of the transaction as registered by the Transaction Tracker's registration daemon.

TYPE: Retained - Character, length 60.

SEE ALSO: Not applicable

ELEMENT DERIVATION: TTANAME is set to the value of the MeasureWare Transaction metric TT_NAME.

RANGE OF VALUE: TTANAME can contain any combination of 1 to 20 alphanumeric characters.

SUMMARIZATION PROCESS: TTANAME is not summarized. TTANAME is used to determine a unique key for records within a file based on the timespan, so other elements in the file can be summarized.

USAGE NOTES: None

T T A P C B X X | Percent Transaction in Bin x

The TTAPCBXx data elements contain the percentage of transactions that completed in the user-defined range during the measurement interval. The first element with this alias is TTAPCB01. The last element with this alias is TTAPCB10.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: TTATRNx - Completed Transactions in Bin x
TTATRANS - Completed Transaction Count

ELEMENT DERIVATION : TTAPCBXx is computed using the following algorithm:

```
TTAPCBx = TTATRANS;
IF TTATRANS
  THEN TTAPCBx = 100*TTATRNx/TTATRANS;
```

RANGE OF VALUE: TTAPCBXx takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: TTAPCBXx contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

T T A P C V S L | Percent Transactions Violated SLO

TTAPCVSL contains percentage of transactions that violated the service level objective set by your installation during the measurement interval.

TYPE: Computed using specific EXPression statements - Numeric, measurement unit is real number.

SEE ALSO: TTAVSLCT - Violated SLO Transaction Count
TTATRANS - Completed Transaction Count

ELEMENT DERIVATION: TTAPCVSL is computed using the following algorithm:

```
TTAPCVSL = TTATRANS;
IF TTATRANS
  THEN TTAPCVSL = 100*TTAVSLCT/TTATRANS;
```

RANGE OF VALUE: TTAPCVSL takes any numeric value ranging from 0 to 100.

SUMMARIZATION PROCESS: TTAPCVSL contains the result of the computation described in ELEMENT DERIVATION. The value is recomputed at each level of summarization.

USAGE NOTES: None

```
+-----------------+
| T T A S L O M X |  Transaction Service Level Objective
+-----------------+
TTASLOMX contains the upper limit on transaction times for
this service level objective.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  TTASLOMX is set to the value of the
MeasureWare Transaction metric TT_SLO_THRESHOLD.

RANGE OF VALUE:  TTASLOMX takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTASLOMX contains the last value of
TTASLOMX processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A T F A I L |  Number of Failed Transactions
+-----------------+
TTATFAIL contains the number of failed transactions during
the last measurement interval for this transaction name.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  TTATFAIL is set to the value of the
MeasureWare Transaction metric TT_FAILED.

RANGE OF VALUE:  TTATFAIL ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTATFAIL contains the last value of
TTATFAIL processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A T R A N S |  Completed Transaction Count
+-----------------+
TTATRANS contains the number of transactions that completed
during the measurement interval.

TYPE:  Accumulated - Numeric, measurement unit is integer.

SEE ALSO:  TTAVSLCT - Violated SLO Transaction Count

ELEMENT DERIVATION:  TTATRANS is set to the value of the
MeasureWare Transaction metric TT_COUNT.

RANGE OF VALUE:  TTATRANS takes any numeric value ranging
from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTATRANS contains the sum of all
values of TTATRANS processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A T R N I D |  Transaction Class ID
+-----------------+
TTATRNID contains the transaction class identifier in this
correlator.

TYPE:  Retained - Numeric, measurement unit is integer.

SEE ALSO:  Not applicable

ELEMENT DERIVATION:  TTATRNID is set to the value of the
MeasureWare Transaction metric TT_CLIENT_TRAN_ID.

RANGE OF VALUE:  TTATRNID ranges from a minimum of 0 to a
maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS:  TTATRNID contains the last value of
TTATRNID processed in the level of summarization.

USAGE NOTES:  None


+-----------------+
| T T A T R N T M |  Cumulative Transaction Response Time
+-----------------+
TTATRNTM contains the cumulative wall clock time for this
transaction during the measurement interval.
```

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: TTATRANS - Completed Transaction Count

ELEMENT DERIVATION: TTATRNTM is computed using MeasureWare metrics in the following calculation:

```
TT_WALL_TIME_PER_TRAN*TT_COUNT
```

RANGE OF VALUE: TTATRNTM takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: TTATRNTM contains the sum of all values of TTATRNTM processed in the level of summarization.

USAGE NOTES: None

T T A T R N X X | Completed Transactions in Bin x

The TTATRNxx data elements contain the number of transactions that completed in this range during the measurement interval. The first element with this alias is TTATRN01. The last element with this alias is TTATRN10.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: TTATRNxx is set to the value of the MeasureWare Transaction metric TTBIN_TRANS_COUNT_xx.

RANGE OF VALUE: TTATRNxx takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: TTATRNxx contains the sum of all values of TTATRNxx processed in the level of summarization.

USAGE NOTES: None

T T A U P R X X | ALIAS TTAUPRXx

The TTAUPRxx elements contain the highest transaction time for the Transaction Tracker bin. The first element with this alias is TTAUPR01. The last element with this alias is TTAUPR10.

TYPE: Retained - Numeric, measurement unit is integer.

SEE ALSO: Not applicable

ELEMENT DERIVATION: TTAUPRxx is set to the value of the MeasureWare Transaction metric TTBIN_UPPER_RANGE_x.

RANGE OF VALUE: TTAUPRxx takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: TTAUPRxx contains the last value of TTAUPRxx processed in the level of summarization.

USAGE NOTES: None

T T A V S L C T | Violated SLO Transaction Count

TTAVSLCT contains the number of transactions that exceeded the service level objective's threshold during the measurement interval.

TYPE: Accumulated - Numeric, measurement unit is integer.

SEE ALSO: TTATRANS - Completed Transaction Count

ELEMENT DERIVATION: TTAVSLCT is set to the value of the MeasureWare Transaction metric TT_SLO_COUNT.

RANGE OF VALUE: TTAVSLCT takes any numeric value ranging from 0 to a maximum that is limited only by practical bounds.

SUMMARIZATION PROCESS: TTAVSLCT contains the sum of all values of TTAVSLCT processed in the level of summarization.

USAGE NOTES: None