IBM Measured Usage Joint Study Technical Customization Document

Document Number CIER-MULCDOC-01

May 26, 1998

IBM Corporation

Poughkeepsie, N.Y.

Contents

Figures

Trademarks

The following terms denoted by an asterisk this publication are trademarks of the IBM Corporation in the United States or other countries or both:

- CICS
- CICS/ESA
- CICS/MVS
- DATABASE 2
- DB2
- IBM
- IMS/ESA
- ISPF/PDF
- MQSeries
- MVS/ESA
- OS/390
- PR/SM
- RMF

About This Document

This document provides a detailed technical description of the MVS/ESA and OS/390 system customization and subsystem product customization required for participation in the Measured Usage Joint Study, which IBM is conducting with MVS/ESA and OS/390 customers. This document lists the specific products and product versions eligible to participate in the joint study. The logistics of collecting data and sending it to IBM for analysis are also covered.

Document's Audience

This document is intended for use by IBM field reps to understand the technical details of the joint study's customization requirements.

This document is also intended for MVS/ESA and OS/390 customers who are considering participating in the Measured Usage Joint Study. Participation in the joint study requires some degree of customization on the customer's part, with the degree of customization dependent on how closely the customer's MVS/ESA and OS/390 systems already conform to the study's minimal requirements. Since this may mean significant customization on the part of some customers, they need to use this document to understand the extent of the customization before deciding whether or not to participate. If they do decide to participate, this document serves as a guide to the customization process with suggestions, examples and checklists.

All customer questions about the joint study in general and about the procedures and guidelines contained in this document should be directed to the appropriate IBM representative. See Appendix A, "Contacting IBM Joint Study Support in the US" on page 27 for details.

Related Information

This document references information in other IBM publications using shortened versions of the publication's title. The following table shows the shortened titles and the complete titles of the publications you might need while you are using this document. The order numbers are listed only for those publications which are MVS/ESA and OS/390 version and release independent.

Short title used in this document	Title
DFP Access Method Services	MVS/DFP: Access Method Services for the Integrated Catalog Facility
DFP Utilites	MVS/DFP: Utilities
JCL Guide	MVS/ESA JCL Guide
JCL Ref	MVS/ESA JCL Reference
MVS Init and Tuning Guide	MVS/ESA Initialization and Tuning Guide
MVS Init and Tuning Ref	MVS/ESA Initialization and Tuning Reference
MVS System Commands	MVS/ESA System Commands
RMF User's Guide	MVS/ESA Resource Measurement Facility User's Guide
System Management Facilities	MVS/ESA System Management Facilities

IBM Measured Usage Joint Study Customization

Overview

The objective of this joint study is to assess the current world wide usage of specific MVS/ESA and OS/390 subsystems. The subsystems participating in this study are limited to: CICS, DB2, IMS and MQSeries.

This study will take advantage of existing MVS/ESA and OS/390 mechanisms to collect system and subsystem usage statistics. Customers may be required to customize their systems and must provide information on their system configurations. RMF is a pre-requisite product or pre-requiste OS/390 feature for the joint study. It will be used to collect a portion of the usage statistics and write them to SMF for the entire length of the study. The required joint study customization consists of:

1. insuring that MVS/ESA and the specified subsystem products have the specified maintenance applied prior to starting data collection
2. generating RMF type 70 records synchronized to the hour
3. enabling and collecting SMF type 0, 70 and 89-1 records

Although the current IBM Measured Usage License Charging (MULC) uses the same SMF type 89-1 records as those collected in this study, the study's collection of these records will have no effect on a customer's current use or non-use of MULC. Customers who use MULC need to continue to collect the required SMF type 89-1 records and provide IBM with required "Software Usage Report" on their existing measurement intervals.

While potentially requiring some customization on the customer's part, this joint study attempts to minimize such changes by attempting to co-exist with and use as many of the customer's current system parameters and RMF options as is possible.

For planning purposes, the joint study estimates that the system overhead of adding SMF type 89 record collection for CICS, IMS and MQSeries will consume less than 1% of those systems with SMF collection already active. The overhead to collect SMF type 89 records for DB2 may be slightly larger than 1% of the total DB2 usage depending on the specific DB2 functions being used.

Each customer needs to assess the impact of participating in this joint study on their operations and system performance.

IBM Measured Usage Joint Study Pre-Reqs

In order to participate in the IBM Measured Usage Joint Study, customers:

1. must have MVS/ESA SP5 or OS/390 and the RMF version/release corresponding to the MVS/ESA or OS/390 release
2. must have at least one of the qualifying product versions listed in Figure 1
3. must be able to send the usage data to IBM on 3480 or 3490 tape cartridges.

All qualifying subsystem products need not be present on a processor for that processor to participate in the joint study. However, all subsystem products, which are present on the processor, must participate in the joint study with the possible exception of DB2. See "DB2" on page 13 for details. In order for a processor to participate, all MVS/ESA images, starting with MVS/ESA SP 3.1.3, and all OS/390 images on that processor footprint, including test LPARs must participate. The study will collect usage data for subsystem versions which are earlier than the qualifying versions listed in Figure 1. See the specific subsystem section of this document for further product details, including required service levels for MVS/ESA, OS/390 and the subsystem products.

Figure 1. Qualifying Subsystem Products		
Subsystem	**Product Name**	**Product Number**
CICS	CICS/ESA Version 4	5655-108
	CICS TS Version 1	5655-147
DB2	IBM DB2 Version 3	5685-DB2
	IBM DB2 for MVS/ESA Version 4	5695-DB2
	IBM DB2 for MVS/ESA Version 5	5655-DB2
IMS	IMS/ESA Version 5	5695-176
	IMS/ESA Version 6	5655-158
MQSeries	MQSeries for MVS/ESA Version 1	5695-137

Customer Obligations for Joint Study Participation

The customer agrees to make the changes described in this document to their participating MVS/ESA (SP3.1.3 or later) and OS/390 systems to collect the required usage data 24 hours a day, for the 30 contiguous days of the study, planned and unplanned system outages, excepted. The customer must provide *Configuration Statements* as described in "Configuration Statements" on page 16 and start data collection on all participating MVS/ESA and OS/390 systems not later than June 14, 1998. At the conclusion of the joint study, the *Configuration Statements* together with all the data collected during the study should be sent together on 3480 or 3490 tape cartridges to IBM via 2 day express delivery. A customization checklist is provided in Appendix B, "Joint Study Customization Check List" on page 32 for your convenience.

Note: All of the customization actions described in this document must be implemented on all participating MVS/ESA and OS/390 systems on a processor footprint before that footprint is considered participating in the study.

Processor Footprint Requirements

Since one of the objectives of the study is to understand subsystem usage by processor footprint, all the MVS/ESA and OS/390 systems running under PR/SM LPARs or VM (hereafter, referred to collectively as MVS/ESA and OS/390 images) on that processor footprint, including test LPARs, must participate in the study.

Whenever a qualifying version of a product executes on any MVS/ESA or OS/390 system on a participating processor footprint, all versions of that product, including earlier versions, must be customized on all images on that processor footprint. See each of the subsystem product sections for a list of the qualifying product versions and required customization.

The study must receive the **same** 30 days of data from all MVS/ESA and OS/390 systems on the participating footprint(s). Although the customization and data collection can start on different days on each of the MVS/ESA and OS/390 images, the study doesn't begin until **all** of the MVS/ESA images on a particular processor footprint are collecting usage data starting at midnight. See the example in "Procedures for Sending SMF Data to IBM" on page 24 for instructions on how to send data for only the days in the study.

MVS/ESA, OS/390, SMF Customization

MVS/ESA and OS/390

MVS/ESA and OS/390 Product Versions

In order for a processor to participate in the joint study, at least one executing system image must be an MVS/ESA SP5 or OS/390 system **and** must have at least one qualifying subsystem version as defined in Figure 1 on page 1. Once a processor meets those requirements, all MVS/ESA system images, starting with MVS/ESA SP3.1.3, and all OS/390 system images on that processor must participate.

MVS/ESA and OS/390 Service Pre-reqs

Figure 2 lists both required and optional service specific to MVS/ESA and OS/390 for participation in the joint study. Note that the required service is fairly old and likely already installed on most customer systems.

			Releases Affected			Available
			MVS/ESA		**OS/390**	
R/O	**Apar #**	**Description**	3 4 4 4 5 5 5 1 2 2 3 1 2 2 3 0 2 0 0 0 2		1 1 1 1 1 1 2 3 4 5	**Available**
R	OW02855	Original SMF type 89 support	X X X X X			05/27/94
R	OW05616	SMF type 89 support	X X X X X			06/27/94
R	OW07260	SMF type 89 support	X X X X X			09/21/94
R	OW08302	SMF type 89 support	X X X X X			10/27/94
R	OW08665	SMF type 89 support	X X X X X			12/18/94
R	OW09370	SMF type 89 support	X X X X X			12/18/94
R	OW10259	DB2 Performance Path Problem	X X X X X			02/14/95
R	OW10721	ABEND0C9 IN IEEMB836	X X X X X			03/21/95
R	OW11636	LOOP IN IGX00025	X X X X X			03/29/95
O	OW16176	DB2 Queries from CICS	X X X		X	04/08/96
O	OW16529	MSGIFA322E MSGIFA343E IN MULC	X X X		X X	05/13/96
O	OW29266	MSGIFA322E MSGIFA343E IN MULC	X			11/12/97
O	OW30153	IMS Queries from CICS	X X X		X X X X X	03/06/98

Figure 2. MVS/ESA and OS/390 Service for Joint Study.

R - Required for participation in the joint study.
O - Optional for participation in the joint study.

Although these are the only MVS/ESA and OS/390 APARs required to participate in the study, IBM recommends that you always apply the latest maintenance for all products.

IPL Requirement

An IPL will be required if any service listed in Figure 2 is applied prior to beginning the joint study. An IPL may also be required if any of the subsystem service listed in "Subsystem Customization" on page 13 is applied.

SMF collection

1. The customer must collect SMF record Types 0, 70 and 89-1 from each LPAR on a participating processor over the course of the **entire** joint study period. The Type 89-1 records must be collected from both the SYS and SUBSYS options. Many customers already meet all of these requirements. Additional record types can be collected at the customer's discretion but are not required by this joint study. The collection of SMF records is controlled by member SMFPRMxx in SYS1.PARMLIB. The SMFPRMxx member is described in *MVS Init and Tuning Ref* and *System Management Facilities*.

   Most customers have daily SMF dump jobs whose "Summary Activity Report", can be used to verify that SMF and RMF are producing the records required by the study. Figure 5 on page 9 contains a portion of a sample "Summary Activity Report". This report only lists the SMF records that have been written for the specified interval. The line with **A** indicates that a Type 0 record was written. This will only occur as a result of an IPL and so will not appear on all reports. The lines with **B** should always be present, since these records should be collected for the study. Note that the appearance of Type 89 records on this report is no guarantee that Type 89 records are being collected from both the SYS and SUBSYS options.

2. Though the customer is free to use SMF interval accounting, this option is not required by the study.

3. Customers using the INTVAL and SYNCVAL parameters (supported on MVS/ESA SP4.3 and higher releases) must meet two conditions:

   a. An INTVAL value that divides evenly into 60
   b. A SYNCVAL of 0 (or a SYNCVAL of 59 for MICS users)

   If the customer also uses the INTERVAL parameter with the INTVAL and SYNCVAL parameters, the study recommends that the INTERVAL parameter either use the SMF option or use a value that divides evenly into the INTVAL parameter. Each of the following three examples meet all the requirements.

   Example 1:
   ```
   INTVAL(60)
   SYNCVAL(00)
   INTERVAL(SMF,SYNC)
   ```

   Example 2:
   ```
   INTVAL(60)
   SYNCVAL(00)
   INTERVAL(006000)
   ```

   Example 3:
   ```
   INTVAL(30)
   SYNCVAL(59)
   INTERVAL(SMF,SYNC)
   ```

4. New SMFPRMxx options can be instituted via an MVS SET command (i.e. `SET SMF=xx` where xx correspond to the suffix for the SMFPRMxx member) without the need for an MVS IPL.

5. The customer must make a good faith effort to save all the SMF data accumulated over the study and must be willing to provide it to IBM for analysis. See "Sending Joint Study Data to IBM" on page 22 before sending SMF data to IBM.

6. Since the customer has access to the original data, the SMF data provided to IBM will not be returned to the customer.

7. The customer is prohibited from deleting or modifying the SMF type 0, 70 and 89-1 records, either through MVS system exits, RMF user exits or other means.

Valid SMFPRMxx Options

Example 1: Minimum SMFPRMxx Options

The following sample includes an example of all the SMF options needed to enable the minimal SMF recording required for the joint study. Note that OS/390 introduced SMF Type 89 subtype 2 records. Those subtype 2 records are not required by the study, but will be collected with the Type 89 subtype 1 records unless "89(1)" is specified. ***Customers not collecting SMF records prior to this study should consider using this sample to collect the minimum set of SMF records.***

```
ACTIVE                      /* ACTIVE SMF RECORDING          */
DSNAME(SYS1.MANX,SYS1.MANY) /* SMF DATASETS                  */
NOPROMPT                    /* DO NOT PROMPT THE OPERATOR    */
MAXDORM(3000)               /* WRITE AN IDLE BUFFER AFTER 30 MIN*/
SID(PROD)                   /* SYSTEM ID FOR PROD SYSTEM     */
LISTDSN                     /* LIST DATA SET STATUS AT IPL   */
SYS(TYPE(0,70,89(1)))       /* COLLECT ONLY JOINT STUDY RECORDS */
```

The following examples will focus only on the SYS and SUBSYS parameters of the SMFPRMxx member, since those are the most critical parameters.

Example 2: Coexistence with existing SYS Options

The following simple example shows existing SYS options which are already in use by the customer and which will collect the SMF record types required by the study, as well as, additional record types deemed important by the customer. ***Customers not collecting SMF records prior to this study should not use this example due to the potentially large volume of SMF records produced.***

```
ACTIVE                      /* ACTIVE SMF RECORDING          */
...
SYS(TYPE(0:255))            /* WRITE ALL RECORDS AS DEFAULT */
```

Example 3: Coexistence with existing SYS, SUBSYS and NOTYPE Options

The following more complex SYS and SUBSYS example uses the NOTYPE parameter and shows existing options are already in use by the customer. These options will collect the SMF record types required by the study, as well as, additional record types deemed important by the customer.

```
ACTIVE                      /* ACTIVE SMF RECORDING          */
...
SYS(NOTYPE(14,15,40,60,62,64),
 EXITS(IEFU83,IEFU84,IEFACTRT,IEFUJV,
       IEFUSI,IEFUTL,IEFU29),NOINTERVAL,DETAIL)

/* WRITE ALL RECORDS AS THE SYSTEM DEFAULT EXCEPT FOR
   THOSE SPECIFICALLY LISTED.  TAKE ALL KNOWN EXITS
   NOTE: JES EXITS CONTROLED BY JES. */

SUBSYS(STC,NOTYPE(14,15,40,57,60,62,64),
     EXITS(IEFU29,IEFU83,IEFU84))

/* WRITE ALL RECORDS AS THE SYSTEM DEFAULT EXCEPT FOR
   THOSE SPECIFICALLY LISTED.  TAKE ONLY THREE EXITS.
   NOTE: IEFU29 EXECUTES IN THE MASTER ASID WHICH IS A
   STC ADDRESS SPACE SO IEFU29 MUST BE ON FOR STC.
   USE ALL OTHER SYS PARAMETERS AS A DEFAULT  */
```

Invalid SMFPRMxx SYS Options

Example 4: **Invalid** SYS Options

The following simple example shows an **incorrect** set of SYS options. This example **does not** provide for the collection of type 89 records.

```
ACTIVE                       /* ACTIVE SMF RECORDING         */
...
SYS(TYPE(0:70))
```

Example 5: **Invalid** SYS Options

The following complex example also shows an **incorrect** set of SYS options. This example **does not** provide for the collection of type 89 records for the system.

```
ACTIVE                       /* ACTIVE SMF RECORDING         */
...
SYS(NOTYPE(14,15,40,60,62,64,82:255),
 EXITS(IEFU83,IEFU84,IEFACTRT,IEFUJV,
       IEFUSI,IEFUTL,IEFU29),NOINTERVAL,DETAIL)

/* WRITE ALL RECORDS AS THE SYSTEM DEFAULT EXCEPT FOR
   THOSE SPECIFICALLY LISTED.  TAKE ALL KNOWN EXITS
   NOTE: JES EXITS CONTROLED BY JES. */

SUBSYS(STC,NOTYPE(14,15,40,57,60,62,64),
    EXITS(IEFU29,IEFU83,IEFU84))

/* WRITE ALL RECORDS AS THE SYSTEM DEFAULT EXCEPT FOR
   THOSE SPECIFICALLY LISTED.  TAKE ONLY THREE EXITS.
   NOTE: IEFU29 EXECUTES IN THE MASTER ASID WHICH IS A
   STC ADDRESS SPACE SO IEFU29 MUST BE ON FOR STC.
   USE ALL OTHER SYS PARAMETERS AS A DEFAULT  */
```

Suggested Changes to SMF Dump Jobs

Although the customer is free to choose both the frequency with which the SMF data is dumped and the names used for saving the SMF files prior to placing them on the tape to be sent to IBM, this document will make recommendations in both areas. The sample JCL provided in this document will support the dumping and saving of the SMF data on a daily basis so that the customer's existing dumping jobs are impacted least.

To facilitate the separation of the study's SMF data from the customer's own SMF data requirements, the customer's existing SMF dump job(s) can be modified to write the joint study's usage data directly to a preallocated joint study dataset, which is unique to the system being measured, as well as, to the customer's existing dataset(s). Writing the study's SMF records to a special usage dataset immediately after the SMF data is being dumped to the customer's usual dump dataset ensures that the usage data is saved through the entire study and eliminates the need to coordinate the activities of multiple MVS/ESA and OS/390 systems. Figure 3 on page 7 contains the sample JCL to preallocate a different usage dataset to hold the joint study's SMF data for each participating MVS/ESA or OS/390 system. Each of these joint study datasets will include the four character MVS system id, as defined by the SID parameter in the SMFPRMxx member of SYS1.PARMLIB, as one of the dataset's qualifying names. The space requirements in this example should be changed to reflect those of the data center. The suggested summary dataset naming convention is:

   hlq.USAGE.*sysid*.SMFDATA

where:

> *hlq*   corresponds to a high level qualifier chosen to conform to data center standards.

sysid   corresponds to the MVS system id from which the data was collected. This is the same as the SID parameter defined in the SMFPRMxx member of SYS1.PARMLIB. See *MVS Init and Tuning Ref* for further details.

```
//........ JOB  ..........
//ALLOC1   EXEC PGM=IEFBR14
//DB2SUM   DD  DSN=hlq.USAGE.sysid.SMFDATA,DISP=(NEW,CATLG),
//             UNIT=SYSDA,SPACE=(CYL,(75,25)),
//             DCB=(RECFM=VBS,BLKSIZE=32760,LRECL=32760)
```

Figure 3. Sample job to allocate SMF usage dataset for MVS/ESA system: *sysid*.

The space requirements for each MVS/ESA and OS/390 system can be approximated by determining the number of SMF Type 0, 70 and 89-1 records typically collected over a 30 day period on each MVS/ESA and OS/390 system. This information can be obtained from the SYSPRINT listing of the jobs used to dump the active SMF datasets on each of the MVS/ESA and OS/390 systems. Multiply the number of SMF type 70 records calculated for 30 days for each system by the average record length for a type 70 record (obtained from same SMF dump job). Repeat the process for the type 89 subtype 1 records and add both results together. This sum should represent the SMF usage requirements for the 30 day study on each system. Use this value to calculate the primary space allocation and use 10-50% of that value to calculate the secondary space allocation.

Be aware that the suggested usage datasets are deliberately allocated large enough to hold all the required data for the joint study. Data will constantly be appended to the end of these datasets, throughout the study. These datasets will, therefore, appear to have too much space allocated at the start of the study and may have their 'free space' reclaimed by any of several storage management products. This should not be permitted to happen, since they may run out of usable extents (space) prior to the end of the study, thereby, resulting in the loss of usage data. This is especially true of the usage study's suggested SMF dataset(s).

Note that suggested procedures and sample JCL, including the space allocations for datasets, presented in this document are guidelines only. It is the customer's responsibility to determine the guidelines' applicability to their data center operations.

Once these study datasets are allocated for each participating MVS/ESA and OS/390 image, the sample LASTSTEP JCL in Figure 4 on page 8 can be added to the customer's existing SMF dump job to continuously add that day's joint study data to the previous day's data, provided the original customer dump step doesn't exclude the study's SMF records from being saved. The example uses *.stepname.ddname to refer to the output file created by the original SMF dump step. If the customer's dump job uses a PROC to invoke the dump program, the backward reference should be of the form *.stepname.procstepname.ddname. See *JCL Ref* for additional information on backward references. The JCL in Figure 4 on page 8 is added as the last step to the customer's existing daily SMF dump jobs, and should not interfere with the existing SMF dump steps. Note that failure to allocate study datasets large enough to hold the specified SMF record types for the entire study or failure to prevent storage management tools from reclaiming the space can cause this job to fail. While the customer's data is correctly dumped, usage data is not saved on the usage dataset. Always check the dump step's return codes to insure that the step ran OK.

```
//....    JOB  ....
//*---------------------------------------------------------
//*  CUSTOMER'S EXISTING SMF DUMP JCL
//*---------------------------------------------------------
//STEP1    EXEC PGM=IFASMFDP
//SYSPRINT DD   ....
//DUMPIN   DD   ....
//DUMPOUT  DD   DSN=customer.dump.dataset,DISP=...
//SYSIN    DD   ....
 ...       ..   ....
//*---------------------------------------------------------
//*  NEW JOINT STUDY DUMP STEP
//*---------------------------------------------------------
//LASTSTEP EXEC PGM=IFASMFDP
//SYSPRINT DD   SYSOUT=A
//DUMPIN   DD   DSN=*.STEP1.DUMPOUT,DISP=SHR
//USAGE    DD   DSN=hlq.USAGE.sysid.SMFDATA,DISP=(MOD,KEEP)
//SYSIN    DD   *
     INDD(DUMPIN,OPTIONS(DUMP))
     OUTDD(USAGE,TYPE(0,70,89(1)))
```

Figure 4. Sample job to write SMF usage data to separate dataset AFTER normal SMF dumping.

Customers who keep all their SMF records for at least 32 days, can avoid changing their existing SMF dump jobs, if they are willing to create a job to extract the study's SMF data from the 30 daily SMF datasets saved during the study and place the data in a single file on the tape cartridge being sent to IBM. Note that the SMF data must be extracted for each participating MVS/ESA and OS/390 image. The JCL for such a job can be modeled on the JCL in Figure 15 on page 24 with the appropriate dump dataset names being used for each of the DUMPIN DD statements.

```
                            SUMMARY ACTIVITY REPORT

     START DATE-TIME  05/11/1998-00:01:03                    END DATE-TIME  05/12/1998-00:00:12

       RECORD        RECORDS       PERCENT     AVG. RECORD    MIN. RECORD    MAX. RECORD      RECORDS
         TYPE           READ      OF TOTAL         LENGTH         LENGTH         LENGTH      WRITTEN

    [A]    0             1         .00 %          60.00             60             60            1

           2             0         .00 %          18.00             18             18            1
           3             0         .00 %          18.00             18             18            1

           8             1         .00 %       5,032.00          5,032          5,032            1
           9             2         .00 %         112.00             60            164            2

          14        12,440        4.01 %         338.50            296            678       12,440
          15         1,196         .39 %         317.62            296            346        1,196

          64         5,564        1.79 %       2,256.48            446          7,050        5,564

    [B]   70            24         .01 %       1,972.00          1,972          1,972           24
          71            24         .01 %       1,449.17          1,308          1,500           24
          72         3,090        1.00 %         900.66            804          7,460        3,090
          73            24         .01 %       5,368.00          5,368          5,368           24
          74           399         .13 %      20,629.33            412         32,744          399
          75           238         .08 %         264.00            264            264          238

          77            24         .01 %       7,760.00            800         32,640           24
          78            48         .02 %       3,619.88          2,448          4,800           48
          79        29,234        9.42 %      17,848.98            346         32,704       29,234
          80           991         .32 %         216.72            180            306          991
          81             1         .00 %       2,512.00          2,512          2,512            1

          88         4,964        1.60 %         268.18            161            280        4,964
    [B]   89            24         .02 %         604.50            336          1,008           24
          90            12         .00 %         330.50             72            628           12

          ..            ..           ..             ..             ..             ..           ..

         150             2         .00 %       1,036.00          1,036          1,036            2
         151             2         .00 %         300.00            300            300            2

     TOTAL        310,339         100 %       3,221.68             18         32,750      310,339


     NUMBER OF RECORDS IN ERROR              0
```

Figure 5. Sample listing from IFASMFDP.

RMF Customization

RMF Product Versions

The joint study requires the RMF version/release appropriate to the level of each of the individual MVS/ESA and OS/390 systems participating in the study.

Processors which do not have an RMF license may participate, but will not be able to supply the RMF portion of the joint study's usage data.

RMF Service Pre-reqs

There is no pre-requisite service required for any of the RMF versions used to produce RMF records for the joint study.

RMF Product Customization

RMF must be customized as described in this section and run on all MVS/ESA and OS/390 system images participating in the study. Many customers already meet both of these requirements.

All the RMF parameters described in this section are described in *RMF User's Guide*.

1. The customer must minimally run an RMF Monitor I Data Gathering session on each participating MVS/ESA and OS/390 system for the entire duration of the joint study.

2. The RMF Monitor I options required by the joint study should be a subset of those options already used by customers running RMF, and should be capable of coexisting with those options. For instance, the study requires that RMF records be written at least once an hour, but shorter intervals are acceptable when those shorter interval values can be divided evenly into 60 so that an interval always ends on the hour. Minimal RMF options include:

CPU	Reports processor activity measurement. This is a required parameter.
INTERVAL	Specifies the length of the reporting interval. Must be chosen so that an integral number of intervals fit within an hour. Values of 5M, 10M, 15M, 20M, 30M and 60M are acceptable for a SYNC value of 0.
NOREPORT	Specifies that no RMF reports are to be generated at the end of the RMF reporting interval. The customer can elect to receive such reports by using the appropriate REPORT value.
NOSTOP	Specifies that the RMF Monitor I session will only be terminated at the request of the operator.
RECORD	Requests that RMF records be written to SMF dataset. This is a required parameter.
SYNC	The joint study requires the synchronization of RMF reporting with each hour. The SYNC parameter specifies whether the RMF interval is synchronized with SMF global sychronization (i.e. `SYNC(SMF)`) or some offset relative to the hour (i.e. `SYNC(RMF,mmM)`). Since RMF recommends synchronizing with the SMF interval, the study will accept `SYNC(SMF)` provided the SMF INTVAL and SMF SYNCVAL options meet the RMF requirements described in this section (i.e. synchronized to the hour). Extreme caution must be exercised if RMF is synchronized with the SMF options, since any changes to those SMF options may result in changes to RMF reporting which may make the RMF data unacceptable to this study.
	The joint study will also accept the `SYNC(RMF,mmM)` form of the SYNC parameter. Specifying a SYNC value of 0M together with any of the acceptable INTERVAL values documented above guarantees that the RMF records are properly synchronized with the hour. A value of 59M will also be accepted, since this is the recommended value for MICS users.

While other values may be acceptable for SYNC, they can often result in situations where RMF reporting does not end on an hour boundary and so should be avoided when possible.

Note: Prior to RMF 4.3.0, RMF's SYNC option syntax was `SYNC(mmM)`. Customers running pre-MVS/ESA SP4.3 systems on a processor participating in the joint study, should be aware of this RMF difference on those systems when customizing RMF for the study.

NOWKLD/WKLD Specifies whether or not workload activity data is to be collected and what type of activity data is to be collected. The study does not require workload activity data so will accept the NOWKLD option, but will also accept the WKLD option with any of its subparameters.

Additional RMF options can be used at the customer's discretion but will not required by this study. A detailed description of all the RMF options can be found in the *RMF User's Guide*.

Valid RMF Options

Example 1: Minimum RMF Options
The following set of options meet the minimal data collection requirements of this study.

```
CPU                     /* COLLECT CPU STATISTICS          */
INTERVAL(60M)           /* SUMMARIZE DATA ONCE EVERY HOUR  */
NOREPORT                /* DON'T PRODUCE REPORT            */
NOSTOP                  /* REMAIN ACTIVE UNTIL OPERATOR STOP */
RECORD                  /* WRITE RECORDS TO SMF            */
SYNC(SMF)               /* SYNC REPORTING WITH SMF         */
NOWKLD                  /* DON'T COLLECT WORKLOAD ACTIVITY */
```

Example 2: Alternative RMF Options
The following set of options also meet the minimal data collection requirements of this study, but produce RMF records every 20 minutes instead of every 60 minutes.

```
CPU                     /* COLLECT CPU STATISTICS          */
INTERVAL(20M)           /* SUMMARIZE DATA EVERY 20 MINUTES */
NOREPORT                /* DON'T PRODUCE REPORT            */
NOSTOP                  /* REMAIN ACTIVE UNTIL OPERATOR STOP */
RECORD                  /* WRITE RECORDS TO SMF            */
SYNC(SMF)               /* SYNC REPORTING WITH SMF         */
NOWKLD                  /* DON'T COLLECT WORKLOAD ACTIVITY */
```

Example 3: Coexistence with existing customer options
The following set of existing customer options is perfectly acceptable, since it also meets the minimal RMF data collection requirements of this study, but captures more data than is required by the study. Note that the interval is synchronized to 1 minute before the hour as recommended for MICS users.

```
CHAN                       /* COLLECT CHANNEL STATISTICS    */
CPU                        /* COLLECT CPU STATISTICS        */
CYCLE(1000)                /* SAMPLE AT ONE SECOND INTERVALS*/
DEVICE(NOSG)               /* NO COLLECT STORAGE GROUP STATS*/
DEVICE(TAPE)               /* COLLECT TAPE DEVICE STATISTICS*/
DEVICE(DASD)               /* COLLECT DASD STATISTICS       */
ENQ(DETAIL)                /* COLLECT ENQ STATISTICS        */
EXITS                      /*         TAKE USER EXITS        */
IOQ(DASD,COMM)             /* COLLECT I/O QUEUING STATS     */
INTERVAL(15M)              /* REPORT AT 15 MINUTE INTERVALS */
NOOPTIONS                  /* OPERATOR MAY EXAMINE AND/OR
                              CHANGE THE RMF OPTIONS        */
NOREPORT                   /* NO WRITTEN REPORTS            */
NOSTOP                     /* REMAIN ACTIVE UNTIL OPER STOP */
PAGESP                     /* COLLECT PAGE SWAP STATISTICS  */
PAGING                     /* COLLECT PAGING STATISTICS     */
RECORD                     /* SMF RECORDING                 */
SYNC(RMF,59M)              /* SYNC FOR MICS                 */
SYSOUT(H)                  /* INTERVAL REPORTS TO CLASS H   */
VSTOR(D,CATALOG,VTAM)      /* COLLECT VIRTUAL STORAGE ACTIV */
WKLD(PERIOD,DOMAIN,GROUP)  /* COLLECT WKLD STATISTICS       */
```

Invalid RMF Options

Example 4: **Invalid** RMF Options
The following set of RMF options are invalid for the joint study for several reasons. The INTERVAL value is
invalid, since the RMF interval will not end on each hour.

```
CPU                        /* COLLECT CPU STATISTICS          */
INTERVAL(18M)              /* SUMMARIZE DATA EVERY 18 MINUTES   */
NOREPORT                   /* DON'T PRODUCE REPORT            */
NOSTOP                     /* REMAIN ACTIVE UNTIL OPERATOR STOP */
RECORD                     /* WRITE RECORDS TO SMF            */
SYNC(SMF)                  /* SYNC REPORTING TO START OF HOUR   */
WKLD(PERIOD,DOMAIN,GROUP)  /* COLLECT WKLD STATISTICS         */
```

Example 5: **Invalid** RMF Options
The following set of RMF options are invalid for the joint study for several reasons. The NORECORD value is
invalid, since none of the RMF statistics will be written to SMF. The SYNC value is invalid since the RMF
interval will not end on the hour based on the INTERVAL value of 30 minutes.

```
CPU                        /* COLLECT CPU STATISTICS          */
INTERVAL(30M)              /* SUMMARIZE DATA EVERY 30 MINUTES   */
NOREPORT                   /* DON'T PRODUCE REPORT            */
NOSTOP                     /* REMAIN ACTIVE UNTIL OPERATOR STOP */
NORECORD                   /* DON'T WRITE RECORDS TO SMF      */
SYNC(RMF,5M)               /* SYNC REPORTING TO START 5 AFTER   */
WKLD(PERIOD,DOMAIN,GROUP)  /* COLLECT WKLD STATISTICS         */
```

Subsystem Customization

CICS

CICS Product Versions
In order for the CICS product to participate, the customer must have at least one copy of CICS/ESA Version 4 or CICS Transaction Server for OS/390 on a participating processor footprint. Once this condition is met, all CICS versions and releases running under all MVS/ESA and OS/390 systems on that participating processor footprint **must** participate in the joint study, including versions or releases no longer supported by IBM (i.e. CICS Versions 2 and 3).

CICS Service Pre-reqs
Figure 6 lists the required service specific to CICS for participation in the joint study. Note that the required service is fairly old and likely already installed on most customer systems.

R/O	Apar #	Description	Affected Releases	Available
R	PN53263	CICS support for SMF type 89s	CICS: 2.1.2, 3.2.1, 3.3	5/27/94
O	PN71234	ABEND0C4 IN DFHJUP	CICS: 2.1.2, 3.2.1, 3.3	5/30/95

Figure 6. CICS Service for Joint Study.

R - Required for participation in the joint study.
O - Optional for participation in the joint study.

DB2

DB2 Product Versions
In order for the DB2 product to participate, the customer must have at least one copy of DB2 Version 3 or Version 4 or Version 5 executing on a participating processor footprint. Once this condition is met, all DB2 Version 2 release 3 and higher DB2 versions running on that participating processor footprint **must** participate in the joint study.

DB2 Service Pre-reqs
Figure 7 lists the required service specific to DB2 for participation in the joint study. Note that the required service is fairly old and likely already installed on most customer systems.

R/O	Apar #	Description	Affected Releases	Available
R	PN63615 PN64972	DB2 support for SMF type 89s	DB2: 2.3, 3.1	1/30/95

Figure 7. DB2 Service for Joint Study.

R - Required for participation in the joint study.
O - Optional for participation in the joint study.

DB2 Customization

If any DB2 subsystem was started while MVS's active SMFPRMxx member was not set to enable the recording of type 89 SMF records for both system and subsystem tasks, those DB2 subsystems must be reinitialized (i.e. be terminated and restarted) after the SMFPRMxx member has been changed to request such type 89 recording and after those changes have been activated with a SET SMF=xx MVS operator command. Failure to reinitialize such DB2 subsystems will result in no SMF data being collected for those subsystems. Customers whose only qualifying subsystem is DB2 and who can not reinitialize all their DB2 subsystems on participating processors, can not participate in the joint study, since no usable usage information will be collected.

If the DB2 system is present on any of the participating MVS/ESA or OS/390 system images, the state of DB2's processor-wide participation must be reflected in columns 60-66 of the appropriate PROCESSOR *Configuration Statement*, as described in "PROCESSOR" on page 18. Note that each participating processor must indicate its DB2 status, even if DB2 is not present or is not reinitialized to collected SMF type 89 records.

IMS

IMS Product Versions

In order for the IMS products to participate, the customer must have at least one copy of IMS/ESA Version 5 or 6 executing on a participating processor footprint. Once this condition is met, all IMS versions and releases running on that participating footprint **must** participate in the joint study, including versions or releases no longer supported by IBM (i.e. IMS DM Versions 3 and 4 and IMS TM Versions 3 and 4).

IMS Service Pre-reqs

Figure 8 lists the required service specific to IMS for participation in the joint study. Note that the required service is fairly old and likely already installed on most customer systems.

R/O	Apar #	Description	Affected Releases	Available
R	PN50024 PN50028	IMS support for SMF type 89s	IMS DM 3.1, IMS TM 3.1	5/27/94
R	PN50029 PN50030	IMS support for SMF type 89s	IMS DM 4.1, IMS TM 4.1	5/27/94
R - Required for participation in the joint study. **O** - Optional for participation in the joint study.				

Figure 8. IMS Service for Joint Study.

MQSeries

MQSeries Product Versions

In order for the MQSeries product to participate, the customer must have at least one copy of MQSeries Version 1 executing under MVS/ESA and OS/390 on a participating processor footprint. Once this condition is met all, MQSeries products running under all MVS/ESA and OS/390 systems on that processor footprint **must** participate in the joint study.

MQSeries Service Pre-reqs

Figure 9 lists the required service specific to MQSeries for participation in the joint study. Note that the required service is fairly old and likely already installed on most customer systems.

R/O	Apar #	Description	Affected Releases	Available
R	PN58961	Support SMF type 89s	MQSeries: 1.1	5/30/95
R	PN73811	High MQSeries usage	MQSeries: 1.1	11/30/95
R	PN80497	High MQSeries usage	MQSeries: 1.1	4/25/96

Figure 9. MQSeries Service for Joint Study.

R - Required for participation in the joint study.
O - Optional for participation in the joint study.

Configuration Statements

Customers are required to create a set of *Configuration Statements* to identify the customer and the customer's hardware and software configurations. This identification is accomplished by creating the control statements described in this section. ISPF can be used to allocate an FB (fixed block) dataset named: *userid*.USAGE.CONFIG having a record length of 80 and a block size of 800 and to create the *Configuration Statements* described below. These *Configuration Statements* will constitute the first file on the 3480/3490 tape cartridges being sent to IBM. *Customers providing usage data from multiple data centers may create a set of Configuration Statements for each data center, when each data center sends its usage data to IBM, separately.*

Some statement fields will require data found on the "IBM Statement of Work for Special Projects - Joint Study" form Z125-4911-00, for "IBM Measured Usage Joint Study", which was provided by IBM at the start of the study.

Order of Configuration Statements

```
*     (comments can be placed anywhere)
CUSTOMER1
CUSTOMER2
CONTACT1
CONTACT2
PROCESSOR
PROCESSOR
 ...
SYSPLEX
SYSPLEX
 ...
```

Statements/Parameters

Comments
Comment statements can be placed anywhere in the control statement input stream.

Column	Contents
1	*
2-71	user comments

CUSTOMER1
The CUSTOMER1 statement identifies the customer by IBM customer number and name. Only one CUSTOMER1 statement is required regardless of the number of participating processors or participating MVS/ESA and OS/390 systems.

Column	Contents
1-9	CUSTOMER1
11-17	IBM Customer number (7 digits with leading zeros, if necessary). This is **not** the enterprise number and **not** the establishment number. This number is listed in the "Signature" Section from IBM's "Statement of Work for Special Projects - Joint Project" form Z125-4911-00. Use the number indicated by **B** in Figure 10 on page 17.
19-48	Customer Name left justified for a maximum of 30 characters. Use the customer name indicated by **A** in Figure 10 on page 17 in the "Signature" Section from IBM's "Statement of Work for Special Projects - Joint Project" form Z125-4911-00.

CUSTOMER2

The CUSTOMER2 statement lists the data center location and country for the customer name listed on the CUSTOMER1 statement. Only one CUSTOMER2 statement is required regardless of the number of participating processors or participating MVS/ESA and OS/390 systems.

Column Contents

1-9   CUSTOMER2

11-40   Customer data center location left justified for a maximum of 30 characters. The customer may use any description they choose to differentiate between customer data centers. This is typically a local address such as a city name.

42-57   Country in which customer data center is located. Please use the exact country spelling as it appears in the following list: Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Israel, Italy, Jordon, Kuwait, Luxembourg, Netherlands, Norway, Poland, Portugal, Russia, Spain, Saudi Arabia, Serbia, Slovak Republic, South Africa, Switzerland, Turkey, United Kingdom, USA.

Figure 10. "Signature" Section from customer's "Statement of Work" form Z125-4911-00

Agreed to: **Customer** name: **Liberty Bell Services** **A**	*Agreed to:* **International Business Machine Corporation**
By _____ **Authorized signature**	By _____ **Authorized signature**
Name (type or print):	Name (type or print): Edward B. Morse
Date:	Date: May 16, 1998
Customer number: **12345** **B**	IBM Customer Agreement Number: .....
Customer Address: 215 Liberty Ave. Philadelphia, PA.	Statement of Work number: IBM Office Address: Somers, New York

---

CONTACT1

The CONTACT1 statement identifies the customer's technical contact who can answer questions which IBM may have about these configuration statements. Only one CONTACT1 statement should be provided regardless of the number of participating processors or participating MVS/ESA and OS/390 systems.

Column Contents

1-8   CONTACT1

10-70   First and last name of the customer's technical contact.

CONTACT2

The CONTACT2 statement provides the phone number and e-mail address of the customer's technical contact. Only one CONTACT2 statement should be provided regardless of the number of participating processors or participating MVS/ESA and OS/390 systems.

Column Contents

1-8   CONTACT2

9-40   Technical contact's full phone number including country code and phone extension (if no direct outside line). North American customers need not provide the country code.

42-74   Technical contact's e-mail address or the value "NONE", if they have no e-mail address.

PROCESSOR

The PROCESSOR statement identifies the processor(s) participating in the study. Each participating processor requires its own PROCESSOR statement. IBM requests that the processor(s), which are listed in the **Products** section of form Z125-4911-00 for the "IBM Measured Usage Joint Study", be used in the study. Some processors listed on this form will be fully identified using processor type, model and serial number (e.g. 9672-R65, 00822). Others will simply be identified by the manufacturer, either with an incomplete model or no other information (e.g. Amdahl). When only the manufacturer is listed, and the customer only has one processor from that manufacturer, there should be no confusion over which processor the joint study is requesting, since only one processor matches that description. If multiple processors fit that description please contact the Joint Study support team for assistance as indicated in Appendix A, "Contacting IBM Joint Study Support in the US" on page 27. If the processor that IBM requested was upgraded or replaced by a different processor, IBM requests that the upgraded or replacement be used in the study and the proper code be entered in column 58 of the PROCESSOR Configuration Statement.

Column Contents

1-9 PROCESSOR

11-20 First 10 digit serial number returned from the MVS "D M=CPU" command issued from any MVS system on the processor. See **C** in Figure 11 on page 19.

22-24 Manufacturer Id field, when present, in the CPU ND result of the MVS "D M=CPU" command issued above. See **D** in Figure 11 on page 19. Leave this field blank if it doesn't appear in field **D**. Note that the CPU Node Descriptor is not supported in IBM processors earlier than the 9021 511 and 711 families.

26-31 The last 6 digits of the CPC ND result, when present, from the MVS "D M=CPU" command issued above. See **E** in Figure 11 on page 19. Leave this field blank if it doesn't appear in field **E**. Note that the CPU Node Descriptor is not supported in IBM processors earlier than the 9021 511 and 711 families.

33-36 The processor's TYPE value, left justified, as found in the appropriate TYPE column of Figure 19 on page 34, based on the manufacturer's name for the processor family. If the processor manufacturer or the manufacturer's TYPE is not listed in Figure 19 on page 34 enter a value of "0000" in this field and place the first 8 characters of the processor manufacturer's name in columns 47-56.

38-45 The processor's MODEL value, left justified, as found in the appropriate MODEL column of Figure 19 on page 34, based on the manufacturer's model name for the processor. If the processor manufacturer is listed in Figure 19 on page 34, but the model is not, place the first 8 characters of the model's numeric designation in columns 38-45.

47-56 The first 8 character's of the processor manufacturer's name, when the processor manufacturer is **not** listed in Figure 19 on page 34, otherwise, leave this field blank.

58-58 When the processor being identified is listed on IBM's "Statement of Work for Special Projects - Joint Project" Z125-4911-00, use the letter immediately following the CPU Serial # as indicated by **F** in Figure 12 on page 19. If the processor is an upgrade or replacement for the processor on form Z125-4911-00, place a "C" in column 58. If the processor being identified by this statement is not listed on IBM form Z125-4911-00, place a "Z" in column 58. This will be the case when the customer chooses to collect data for processor's in addition to or instead of those processors specified by IBM on form Z125-4911-00.

60-66 The Julian date in the format yyyyddd representing the first full day on which all DB2 systems on all LPARs on this processor started collecting data for SMF type 89 records.
- Use the date of the first full day of the study, if all DB2 LPARs on the processor are already collecting data for SMF type 89 records at the start of the study.
- Use a value of "0000000" to indicate that **no** DB2 subsystems are run on any LPAR on this processor during the **entire** joint study.
- Use a value of "9999999" to indicate that at least 1 DB2 subsystem could **not** collect SMF Type 89 records because it required reinitialization and could not be reinitualized before the joint study ended.

```
IEE174I 14.04.53 DISPLAY M 395
PROCESSOR STATUS
ID  CPU              SERIAL
0    +               0108229672  C
1    +               1108229672
2    +               2108229672
3    +               3108229672
4    +               4108229672
5    +               5108229672
CPC ND = 009672.R65.IBM.02.000000000822
                    D                E

CPC ID = 00
+ ONLINE    - OFFLINE    . DOES NOT EXIST
CPC ND  CENTRAL PROCESSING COMPLEX NODE DESCRIPTOR
CPC ID  CENTRAL PROCESSING COMPLEX IDENTIFIER
```

Figure 11. Sample listing from D M=CPU.

3. Products

......

```
CPU Serial #        _9672-R65__00822_____    A   F
                    _Amdahl_____     B
                    _____    ___
                    _____    ___
                    _____    ___
                    _____    ___
```

Figure 12. Section 3 from customer's "Statement of Work", IBM form Z125-4911-00.

SYSPLEX

The SYSPLEX statement identifies the sysplex(s), if any, configured on each processor participating in the study, and the associated SYSNAMEs of the systems in that sysplex, regardless of whether those systems are on a participating processor or not. The sysplex name and the SYSNAMEs of all its currently configured systems can be obtained by issuing the MVS "D XCF" command from any system in the sysplex. The resulting IXC334I message lists the sysplex name and the SYSNAMEs of all configured systems for that sysplex. The "D XCF" command must be issued from one system in each SYSPLEX, if systems from multiple sysplexes participate in the joint study. A SYSPLEX statement must be provided for each sysplex name / SYSNAME combination.

No SYSPLEX statement is required if no sysplexes were present on any of the processors participating in the joint study.

Column	Contents
1-7	SYSPLEX
9-16	The 8 character sysplex name from the MVS "D XCF" command, issued from any sysplex system on each participating processor.
18-25	The system's 8 character name from the list of SYSNAMEs returned from the MVS "D XCF" command belonging to the sysplex named in columns 9-16.

Configuration Statements Example

```
*---+----1----+----2----+----3----+----4----+----5----+----6----+----7----
*
CUSTOMER1 0012345 Liberty Bell Services
CUSTOMER2 Philadelphia, Pa              USA
*
CONTACT1 Fred Smith
CONTACT2 (610) 555-2031 EX 3491         fred_smith@nsp.lbells.com
*
*------------------------ PROCESSORS ----------------------------------
*
PROCESSOR 0108229672 IBM 000822 9672 R65               A 1998152
PROCESSOR 0305710500 ADH 000571 0500 GS545E            B 1998172
PROCESSOR 0304989021 HTC 000498 HDS  SKY-11            Z 9999999
PROCESSOR 0232863090          0000 M80/43   MAGNUSON   Z 0000000
*
*-------------------------SYSPLEXS-------------------------------------
*
SYSPLEX REGIONS  REGION1
SYSPLEX REGIONS  REGION2
SYSPLEX REGIONS  REGION3
SYSPLEX REGIONS  REGION4
SYSPLEX INTERNAL BILLING
SYSPLEX INTERNAL RECEIVNG
SYSPLEX INTERNAL MARKETNG
SYSPLEX INTERNAL PAYROLL
```

Figure 13. Sample Configuration Statements for the Libery Bell Services Comany.

The sample configuration statements shown in Figure 13 are for the mythical Liberty Bell Services Company. The company has four processors participating in the joint study. The first two were requested by IBM and the third and fourth processors were requested by the Liberty Bell Services, itself.

The 9672-R65 processor was listed on IBM form Z125-4911-00 (see Figure 12 on page 19) and was followed by the letter "A". Several LPARs on that processor run DB2 subsystems. All of the DB2 subsystems were reinitialized after each corresponding OS/390's SMFPRMxx member was set to collect SMF Type 89 records and before the start of the joint study. As a consequence, the Julian date, on which the processor started the joint study, was placed in columns 60-66.

The IBM form Z125-4911-00 only listed "Amdahl" as the second processor. Since the customer only had one Amdahl processor, a Millennium Global Server 545E, that was the one used in the joint study. Since "Amdahl" was followed by a "B", "B" was placed in column 58. Several LPARs on that processor run DB2 subsystems. Some of the DB2 subsystems could not be reinitialized until after the joint study was underway. As a consequence, the Julian date for the first full day, on which all the DB2 subsystems on all LPARs on that processor were reinitialized, was placed in columns 60-66.

Although not listed on IBM form Z125-4911-00, the customer chose to add their Hitachi Skyline 11 to the study. Since the Skyline was not on the form Z125-4911-00, a "Z" was placed in column 58. One of the DB2 subsystems on the Skyline could not be reinitialized until after the joint study over. As a consequence, a value of "9999999" was placed in columns 60-66.

Although not listed on IBM form Z125-4911-00, the customer also chose to add their Magnuson model M80/43 to the study. Since the MVS "D M=CPU" command did not list any information when issued from a Magnuson

system, columns 22-31 were left blank. Since the Magnuson processor is not listed in Figure 19 on page 34, a value of "0000" was placed in columns 33-36, the manufacturer's model value of "M80/43" was placed in columns 38-43, and the value "MAGNUSON" was placed in columns "47-54". Since the Magnuson was not on the form Z125-4911-00, a "Z" was placed in column 58. Since this processor does not run any DB2 subsystems, a value of "0000000" was placed in columns 60-66.

The customer has two sysplexes, REGIONS and INTERNAL. The REGIONS sysplex has systems with SYSNAMEs of REGION1, REGION2, REGION3 and REGION4. Note that REGION3 and REGION4 reside on a processor which did not participate in the joint study. The INTERNAL sysplex has systems with SYSNAMEs of BILLING, RECEIVNG, MARKETNG and PAYROLL. The customer also had other systems that were not a part of the Parallel Sysplex. Those systems were not listed on any SYSPLEX statement.

Sending Joint Study Data to IBM

The *Configuration Statements* and the SMF data collected during the joint study, must be copied to 3480 or 3490 tape cartridges using the sample JCL provided in this section and sent to IBM immediately after the conclusion of the study for analysis - 31 days after the study started on all LPARs on the processor. Usage datasets from the same processor footprint should be placed on the same tape cartridge, space permitting. The datasets should be placed in the following order on the tape cartridge.

1. File the containing *Configuration Statements*.
2. File containing the SMF data for all MVS/ESA and OS/390 images on all participating footprints.

Customers providing usage data from multiple data centers may create a set of Configuration Statements *for each data center, when each data center sends its usage data to IBM, separately.*

The 3480/3490 tapes, when received, become the property of IBM and will not be returned.

General Instructions

The following steps should be followed when sending the data collected during the joint study to IBM.

1. Send only the data collected from processors participating in the joint study. Customers may have several processor footprints, some of which do not participate in the joint study. Do not send data from the processors which do not participate in the study.

2. The data must be placed on 3480 or 3490 tape cartridges and should use the IDRC (compression) feature, if possible. 3490E tape cartridges with IDRC are preferred. *The study will not accept data sent on any other media.*

3. The 3480/3490 tape cartridges must have standard labels (SL) so that the file DCBs are placed on the internal tape cartridge labels.

   The joint study's SMF data from all participating MVS/ESA and OS/390 systems, should be placed in the same file on the same tape cartridge, space permitting. The customer may create separate SMF files for each processor and place them on the same tape cartridge, if space permits. Sample JCL for copying each of these different types of data to tape cartridges are included in this section.

4. Each tape cartridge **must** be accompanied by a completed *IBM Measured Usage Joint Study Cartridge Form* found on page 26 so that all the files are properly identified.

5. The tape cartridge(s) should be sent using 2 day express mailer to IBM at the following address:

   ```
   IBM CORP.
   SOUTH ROAD
   POUGHKEEPSIE, NEW YORK 12602
   BUILDING 710 LOADING DOCK
   ATTN: JANE GARTLAND
   DEPARTMENT S14A/26
   MAIL STATION: S14A/710-1
   PHONE : (914) 435-3077
   ```

Procedures for Sending Configuration Statements to IBM

The *Configuration Statements*, as defined in "Configuration Statements" on page 16, is the first set of data to be copied to the tape cartridge(s) being sent to IBM. There is only one set of *Configuration Statements* per customer regardless of the number of participating processors. The following sample job should be used to copy the *Configuration Statements* to the joint study tape cartridges being sent to IBM at the conclusion of the joint study. The sample job uses the IEBGENER utility to copy the *Configuration Statements* to the first file on the tape cartridge. IEBGENER is documented in *DFP Utilities*.

```
//....    JOB  ....
//*********************************************
//*       CUSTOMER CONFIGURATION STATEMENTS
//*********************************************
//STEP1    EXEC PGM=IEBGENER
//SYSPRINT DD SYSOUT=A
//SYSUT1   DD DSN=userid.USAGE.CONFIG,DISP=SHR
//SYSUT2   DD DSN=USAGE.CONFIG,UNIT=tapeaddr,
//            DCB=(RECFM=FB,LRECL=80,BLKSIZE=800),
//            DISP=(NEW,KEEP),LABEL=(1,SL)
//SYSIN    DD DUMMY,DCB=BLKSIZE=80
//
```

Figure 14. Sample JCL for copying *Configuration Statements*

Procedures for Sending SMF Data to IBM

The SMF records will comprise the data on the second file of the tape cartridge(s) being sent to IBM. The study prefers that the SMF data from all the participating MVS/ESA and OS/390 images be merged into the second file. If the suggested SMF joint study datasets were created for each participating system per the guidelines in "SMF collection" on page 4, then the sample job in Figure 15 can be used to transfer the all SMF data collected for the joint study to the second file on the tape cartridge. Note that the joint study strongly recommends using the IFASMFDP program to copy the SMF data to the tape cartridge rather than using IEBGENER or any other utility programs. *Note that the following jobs must be run on an MVS/ESA SP5 or OS/390 system, even though some of the SMF data may have been collected on a pre-MVS/ESA SP5 system.*

Before using the sample JCL:

1. Determine the number of systems for which joint study data was collected.
2. Create a //DUMPIN DD statement for each system substituting the appropriate *hlq* and *sysid*n values for each
3. Create a matching INDD SYSIN statement for each //DUMPIN DD statement.
4. Change the *tapeaddr* value to your data center's UNIT value for a 3480 or 3490 tape drive or address.
5. Change the first set of yyyyddd values on the DATE SYSIN statement to the first date on which all images on all processors were collecting SMF data for the study. Change the second set of yyyyddd values to the last date for which SMF data was collected for all processors.

```
//....     JOB  ....
//FILE2    EXEC PGM=IFASMFDP
//SYSPRINT DD   SYSOUT=A
//DUMPIN1  DD   DSN=hlq.USAGE.sysid1.SMFDATA,DISP=SHR
//DUMPIN2  DD   DSN=hlq.USAGE.sysid2.SMFDATA,DISP=SHR
//DUMPIN3  DD   DSN=hlq.USAGE.sysid3.SMFDATA,DISP=SHR
//.......  DD   .......................................
//DUMPINn  DD   DSN=hlq.USAGE.sysidn.SMFDATA,DISP=SHR
//DUMPOUT  DD   DSN=USAGE.SMFDATA,UNIT=tapeaddr,
//              DISP=(NEW,KEEP),LABEL=(2,SL)
//SYSIN    DD   *
     INDD(DUMPIN1,OPTIONS(DUMP))
     INDD(DUMPIN2,OPTIONS(DUMP))
     INDD(DUMPIN3,OPTIONS(DUMP))
     ........................
     INDD(DUMPINn,OPTIONS(DUMP))
     OUTDD(DUMPOUT,TYPE(0,70,89(1)))
     DATE(yyyyddd,yyyyddd)
```

Figure 15. Sample JCL for copying all the SMF data to file 2 of the tape cartridge

If the customer prefers to combine the SMF data from each participating processor on a separate file on the tape cartridge, the sample JCL in Figure 16 on page 25 can be used to create separate files. Note that the joint study strongly recommends using the IFASMFDP program to copy the SMF data to the tape cartridge rather than using IEBGENER or any other utility programs. Also note that the same dataset name is used for all the SMF files on the tape cartridge. Before using the sample JCL:

1. Determine the number of systems for which joint study data was collected for each participating processor.
2. Create a //DUMPIN DD statement for each system on the first processor, substituting the approriate *hlq* and *sysid*n values for each.
3. Create a matching INDD SYSIN statement for each //DUMPIN DD statement.
4. Change the *tapeaddr* value to your data center's UNIT value for a 3480 or 3490 tape drive or address.

5. Change the first set of yyyyddd values on the DATE SYSIN statement to the first date on which all images on that processor were collecting SMF data for the study. Change the second set of yyyyddd values to the last date for which SMF data was collected for all images on that processor.
6. Create a second job step for the second processor following the same directions as the first processor, changing the first value of the LABEL parameter of the //DUMPOUT DD statement to a value one higher than that used in the previous job step.
7. Create additional job steps as needed following the same procedure.

```
//....    JOB  ....
//*-------------------------------------------------------
//*                  PROCESSOR 1
//*-------------------------------------------------------
//FILE2    EXEC PGM=IFASMFDP
//SYSPRINT DD   SYSOUT=A
//DUMPIN1  DD   DSN=hlq.USAGE.sysid1.SMFDATA,DISP=SHR
//DUMPIN2  DD   DSN=hlq.USAGE.sysid2.SMFDATA,DISP=SHR
//DUMPIN3  DD   DSN=hlq.USAGE.sysid3.SMFDATA,DISP=SHR
//.......  DD   .......................................
//DUMPINn  DD   DSN=hlq.USAGE.sysidn.SMFDATA,DISP=SHR
//DUMPOUT  DD   DSN=USAGE.SMFDATA,UNIT=tapeaddr,
//              DISP=(NEW,KEEP),LABEL=(2,SL)
//SYSIN    DD   *
     INDD(DUMPIN1,OPTIONS(DUMP))
     INDD(DUMPIN2,OPTIONS(DUMP))
     INDD(DUMPIN3,OPTIONS(DUMP))

     .........................
     INDD(DUMPINn,OPTIONS(DUMP))
     OUTDD(DUMPOUT,TYPE(0,70,89(1)))
     DATE(yyyyddd,yyyyddd)
//*-------------------------------------------------------
//*                  PROCESSOR 2
//*-------------------------------------------------------
//FILE3    EXEC PGM=IFASMFDP
//SYSPRINT DD   SYSOUT=A
//DUMPIN1  DD   DSN=hlq.USAGE.sysid1.SMFDATA,DISP=SHR
//DUMPIN2  DD   DSN=hlq.USAGE.sysid2.SMFDATA,DISP=SHR
//DUMPIN3  DD   DSN=hlq.USAGE.sysid3.SMFDATA,DISP=SHR
//.......  DD   .......................................
//DUMPINn  DD   DSN=hlq.USAGE.sysidn.SMFDATA,DISP=SHR
//DUMPOUT  DD   DSN=USAGE.SMFDATA,UNIT=tapeaddr,
//              DISP=(NEW,KEEP),LABEL=(3,SL)
//SYSIN    DD   *
     INDD(DUMPIN1,OPTIONS(DUMP))
     INDD(DUMPIN2,OPTIONS(DUMP))
     INDD(DUMPIN3,OPTIONS(DUMP))

     .........................
     INDD(DUMPINn,OPTIONS(DUMP))
     OUTDD(DUMPOUT,TYPE(0,70,89(1)))
     DATE(yyyyddd,yyyyddd)
//*-------------------------------------------------------
//*                  PROCESSOR 3
//*-------------------------------------------------------
//FILE4    EXEC PGM=IFASMFDP
//...
```

Figure 16. Sample JCL for consolidating SMF data on to **multiple** files of the tape cartridge

IBM Measured Usage Joint Study Cartridge Form

Customer name	IBM Customer number (see **B** inFigure 10 on page 17 .)

Data center address (full mailing address)	

Data center operations contact	Phone number: ( )

External 3480/3490 VOLSER	Internal 3480/3490 VOLSER

	Config file	SMF file	Processor Serial Number(s)	Dataset name	RECFM[1]	LRECL[1]	BLKSIZE[1]
File 1							
File 2							
File 3							
File 4							
File 5							
File 6							
File 7							
File 8							

1 - Required only when IFASMFDP was not used to place the file on on tape.

1 of _

Appendix A. Contacting IBM Joint Study Support in the US

The joint study support described in this appendix will be available on business days from 9 AM to 5 PM Eastern Daylight Savings Time. There is no off-shift, weekend or holiday support. The joint study team prefers e-mail questions to phone calls.

The following sections describe how to contact IBM for:

1. Staying current with joint study changes, if they occur.
2. Asking questions about the joint study's Statement of Work and customization.
3. Sending the signed Statement of Work to IBM and requesting a signed copy from IBM.
4. Sending customization changes to IBM for review.
5. Determining MULC prices from joint study data.

Asking Questions About the Joint Study

Customers with questions about the joint study's statement of work or customization can contact the joint study by any of the following means (listed in the joint study's prefered order). Note that sending the IBM Statement of Work back to IBM and sending parmlib members to IBM for review are described in other sections in Appendix A.

Joint Study WEB Site

The joint study is creating a WEB site as the prefered means of communicating with joint study participants. Although the joint study WEB site is an extension of the IBM S/390 Service Update Facility (SUF), the customer need not be a registered SUF user to use the joint study's WEB site. The WEB site's URL is http://www.s390.ibm.com/suf/usage/ and has a target availability date of June 2, 1998. The WEB site can be used to:

1. Submit questions about the joint study
2. View changes to the joint study, if such changes occur
3. View Frequently Asked Questions (FAQs)
4. View instructions on how to determine MULC prices for processors participating in the study
5. View the most current copy of *IBM Measured Usage Joint Study Technical Customization Document* online
6. Download a zipped file containing the the most current BookManager, LIST3820, PostScript and PDF copies of *IBM Measured Usage Joint Study Technical Customization Document*

Joint Study E-Mail Address

Joint study questions can be sent to the joint study's e-mail address: usage@us.ibm.com.

1. Address your e-mail note to: "IBM Usage Study" at address: usage@us.ibm.com.
2. Choose an appropriate category from the following list and place that choice on the SUBJECT line of your e-mail note:
    a. Statement of Work Question
    b. Technical or Customization Question
    c. MULC Related Question
    d. Other Joint Study Question
3. Provide the following information in your e-mail note:
    a. Company Name
    b. City/State Address
    c. IBM customer number (if known)
    d. Your Name
    e. Your complete phone number, including your extension, if you have no direct line.

4. State your joint study question(s)
5. Send the note to IBM.
6. IBM will return review comments to the e-mail address from which the original request was sent.

Joint Study Phone Number

Use the IBM Support Services phone number, 1-800-237-5511, to contact IBM for answers to joint study related questions. *You need not subscribe to the "IBM Support Line" Service to ask a joint study question, as long as that question is routed to the correct "alias", as indicated in step 2.*

1. After reaching IBM Support Services choose Option 4 (OS/390, MVS, VM, VSE), from the automated phone menu.
2. Inform IBM that you have a question to be routed to the "usage study" alias.
3. IBM will request your name, phone number, etc.
4. You will ask IBM your question
5. IBM will call you back with an answer to your question

Returning the Statement of Work to IBM

Before returning the IBM Statement of Work (Z125-4911-00) verify that:

1. The Statement of Work has an authorized signature. (page 1)
2. The signer's name is printed below their signature. (page 1)
3. The Statement of Work is dated. (page 1)
4. CPU Serial #'s are changed or added, if necessary. (page 2)
5. The name and address of the customer's joint study representative is provided. (page 4)

When FAXing the Statement of Work to IBM at 1-914-766-1796

1. Use the "fine" mode on your FAX machine, if available.
2. The joint study prefers that you complete and use the FAX Cover Sheet on Figure 17 on page 30, when FAXing the Statement of Work.
   a. Please enter the date on which you expect to start collecting data for the study, so that the study can notify you when the study is over for that processor(s).
   b. Check off the box on Figure 17 on page 30, if you want IBM to FAX a copy of the completed Statement of Work back to you, and fill out the "return" information on that form.

Sending customization changes to IBM for review

The IBM joint study team will review changes to the customer's ERBRMFxx and SMFPRMxx members of SYS1.PARMLIB, if the customer requests such a review. *Note that the joint study review will only assess whether the revised members meet the joint study's minimal requirements. No exhaustive syntax checking will be performed.*

The joint study team prefers that customers e-mail the members for review but will accept FAXes. Either or both ERBRMFxx and SMFPRMxx parmlib members can be sent for review. Follow one of the following procedures for sending parmlib members to IBM.

Review parmlib changes via E-mail

The following technique requires that your e-mail system must have the capability of attaching files to the e-mail note.

1. Transfer a copy of the revised parmlib member (ERBRMFxx or SMFPRMxx) from your MVS/ESA or OS/390 system to your PC treating your data as TEXT during the transfer process (i.e. convert it to ASCII). Use the member name as the file name and use .REV as the file extension (e.g. ERBRMF22.REV).

2. Transfer a copy of the original parmlib member (ERBRMFxx or SMFPRMxx) that your revised member was based on from your MVS/ESA or OS/390 system to your PC treating your data as TEXT during the transfer process (i.e. convert it to ASCII). Use the original member name as the file name and use .ORI as the file extension (e.g. ERBRMF12.ORI).

3. Repeat the process for the other member to be reviewed, if required.

4. Address your e-mail note to: "IBM Usage Study" at address: usage@us.ibm.com.

5. Use "PARMLIB REVIEW" as the SUBJECT line of your e-mail note.

6. Provide the following information in your e-mail note:

   a. Company Name
   b. City/State Address
   c. IBM customer number (Can be found at **B** in Figure 10 on page 17.)
   d. Your Name
   e. Your complete phone number, including your extension, if you have no direct line.

7. Attach the revised/original set(s) of PC files that were created above to the note

8. Send the note to IBM.

9. IBM will return review comments to the e-mail address from which the original request was sent.

Review parmlib changes via FAX:

1. Print a copy of the revised parmlib member (ERBRMFxx or SMFPRMxx) created for the the study. Write "REVISED" at the top of the listing.

2. Print a copy of the original parmlib member (ERBRMFxx or SMFPRMxx) that your revised member was based on and write "ORIGINAL" at the top of the listing.

3. Repeat the process for the other member to be reviewed, if required.

4. Make a copy of the FAX Cover Sheet from Figure 18 on page 31 and provide the following information as indicate on the Cover Sheet:

   a. Company Name
   b. City/State Address
   c. IBM customer number (Can be found at **B** in Figure 10 on page 17.)
   d. Your Name
   e. Your complete phone number, including your extension, if you have no direct line.
   f. FAX number to return comments to.

5. Fax the cover sheet together with the parmlib listings to IBM at 1-800-319-5777.

6. IBM will return review comments to the FAX number indicated on the FAX Cover Sheet used to provide the data to IBM.

IBM USAGE STUDY - STATEMENT OF WORK

FAX COVER SHEET

Date: _____

From: Company Name: _____

City/State Address: _____

IBM Customer No: _____

Date data collection starting: _____

To: IBM Usage Study: 1-914-766-1796

Total number of pages being transmitted: _____
 (including cover sheet)

☐ Check if completed Statement of Work is to be returned

to: Name: _____

FAX Number: _____

COMMENTS:

_____

_____

_____

_____

_____

_____

Figure 17. FAX Cover Sheet for returning completed Statement of Work (Z125-4911-00) to IBM.

IBM USAGE STUDY - PARMLIB REVIEW

FAX COVER SHEET

Date: _____

From: Company Name: _____
       City/State Address: _____
       IBM Customer No: _____

To: IBM Usage Study: 1-800-319-5777

Total number of pages being transmitted: _____
 (including cover sheet)

Return review comments on parmlib members

to: Name: _____
       Phone Number: _____
       Fax Number: _____

COMMENTS:
_____
_____
_____
_____
_____
_____
_____

Figure 18. FAX Cover Sheet for sending parmlib members to IBM for review.

Appendix B. Joint Study Customization Check List

The following checklist is for each MVS/ESA and OS/390 system on a participating processor. **Note:** All systems must start data collection no later than June 14, 1998.

IPL addr: _____

System id: _____

SYSNAME: _____

Sysplex name: _____

Processor: _____ (type/model/serial)

LPAR name: _____

1. Subsystem customization is complete:

   ___ a. Applied required CICS service to all applicable CICS releases.
   ___ b. Applied optional CICS service, if desired, to all applicable CICS releases.
   ___ c. Applied required DB2 service to all applicable DB2 releases. Otherwise, no DB2 data will be collected.
   ___ d. Applied required IMS service to all applicable IMS releases.
   ___ e. Applied required MQSeries service to all applicable MQSeries releases.

2. RMF customization is complete:

   ___ a. Created appropriate RMF Monitor I options for the joint study.
   ___ b. Instituted a process to insure that RMF is started at IPL with study options.
   ___ c. Reset RMF Monitor I to joint study's options (no IPL needed).

3. MVS system customization is complete:

   ___ a. Applied Required MVS/ESA and OS/390 service.
   ___ b. Applied optional MVS/ESA and OS/390 service, if desired.
   ___ c. Optional:
       ___ 1) Created suggested SMF datasets to hold the study's SMF records as described by "Suggested Changes to SMF Dump Jobs" on page 6.
       ___ 2) Modified and tested existing SMF dump jobs as described by "Suggested Changes to SMF Dump Jobs" on page 6.
       ___ 3) Verified that the study's SMF records are being collected by reviewing "Summary Activity Report" from the daily SMF dump job on this system. See Figure 5 on page 9 for a sample report.
   ___ d. Created an SMFPRMxx member with SMF options required by the joint study.
   ___ e. Ensured that the new SMFPRM member is used at IPL time.
   ___ f. IPLed the system, if MVS/ESA, OS/390, CICS, DB2, IMS or MQSeries service was applied and an IPL was required.
   ___ g. If **no** system IPL was required:
       ___ 1) Issued the MVS SET SMF=xx command to set SMF to the options required by the joint study.
       ___ 2) (optional) Reinitialized (e.g. restart) each DB2 subsystem, if the SMFPRM options had to be changed to produce type 89 records. Otherwise, no DB2 usage data will be collected.

4. Date/Time that the joint study's RMF & SMF collection started for this system.

   Date: _____ (No later than June 14, 1998)

   Time: _____

Joint Study Tape Cartridge Check List

Use the following check list after the joint study data has concluded (31 days after it started on all LPARs on all joint study processors):

___ 1. Created appropriate *Configuration Statements* as described by "Configuration Statements" on page 16 and save *userid*.USAGE.CONFIG as the first file on the tape cartridge being sent to IBM.

___ 2. Moved *userid*.USAGE.CONFIG. to the first file on the tape cartridge being sent to IBM, as described in "Sending Joint Study Data to IBM" on page 22

___ 3. Moved the joint study's SMF data from each participating MVS/ESA and OS/390 system on every participating processor to file 2 (and possibly other files) on the tape cartridge being sent to IBM, as described in "Sending Joint Study Data to IBM" on page 22

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

    Processor: _____   System id: _____

___ 4. Filled out a *IBM Measured Usage Joint Study Cartridge Form* found on page 26 for every tape cartridge being sent to IBM.

___ 5. After all the joint study data has been moved to one or more 3480/3490 tape cartridges, mailed those tape cartridges, together with the *IBM Measured Usage Joint Study Cartridge Form* to IBM via 2 day express deilvery.

Appendix C. Processor Type and Model Values

To select the correct TYPE and MODEL values to place in columns 33-36 and 38-45, respectively, of the PROCESSOR Configuration Statement:

1. Locate the table section for the appropriate processor manufacturer (i.e. IBM, Amdahl, Comparex, Hitachi, National Advanced Systems, Olivetti).
2. Locate the appropriate table entry for your processor
   a. For IBM processors locate your processor using both the TYPE and MODEL values listed in the IBM section of the table.
   b. For all other processors locate your processor using MODEL designation provided by the processor manufacturer.
3. Use the TYPE and MODEL values from the table entry, which was located in Step 2 above, for columns 33-36 and 38-45 of the PROCESSOR Configuration Statement.

Figure 19 (Page 1 of 7). Type and Model values for PROCESSOR Configuration Statement

Type	Model	Type	Model	Type	Model
		IBM Processors			
2003	102	2003	135	2003	216
2003	103	2003	136	2003	224
2003	104	2003	146	2003	225
2003	105	2003	156	2003	227
2003	106	2003	1C5	2003	237
2003	107	2003	203	2003	246
2003	115	2003	204	2003	247
2003	116	2003	205	2003	257
2003	124	2003	206	2003	2C5
2003	125	2003	207		
2003	126	2003	215		
3000	A10	3000	A20		
3090	15T	3090	170S	3090	300S
3090	17T	3090	180E	3090	380J
3090	18T	3090	180J	3090	380S
3090	25T	3090	180S	3090	400E
3090	28T	3090	200E	3090	400J
3090	100S	3090	200J	3090	400S
3090	110J	3090	200S	3090	500E
3090	120E	3090	250J	3090	500J
3090	120J	3090	250S	3090	500S
3090	120S	3090	280E	3090	600E
3090	150E	3090	280J	3090	600J
3090	150J	3090	280S	3090	600S
3090	150S	3090	300E		
3090	170J	3090	300J		
4381	90E	4381	91E	4381	92E

Type	Model	Type	Model	Type	Model
9021	330	9021	720	9021	941
9021	340	9021	740	9021	942
9021	500	9021	820	9021	952
9021	520	9021	821	9021	962
9021	580	9021	822	9021	972
9021	620	9021	831	9021	982
9021	640	9021	832	9021	9X2
9021	660	9021	860		
9021	711	9021	900		
9121	180	9121	440	9121	610
9121	190	9121	480	9121	621
9121	210	9121	490	9121	622
9121	260	9121	511	9121	732
9121	311	9121	521	9121	742
9121	320	9121	522		
9121	411	9121	570		
9221	120	9221	191	9221	221
9221	130	9221	200	9221	421
9221	150	9221	201		
9221	170	9221	211		
9672	R11	9672	R52	9672	R96
9672	R12	9672	R53	9672	RA2
9672	R14	9672	R54	9672	RA4
9672	R15	9672	R55	9672	RA5
9672	R16	9672	R56	9672	RA6
9672	R21	9672	R61	9672	RB4
9672	R22	9672	R63	9672	RB5
9672	R24	9672	R64	9672	RB6
9672	R25	9672	R65	9672	RC4
9672	R26	9672	R66	9672	RC5
9672	R31	9672	R72	9672	RC6
9672	R32	9672	R73	9672	RD6
9672	R34	9672	R74	9672	RX3
9672	R35	9672	R75	9672	RX4
9672	R36	9672	R76	9672	RX5
9672	R41	9672	R83	9672	RX6
9672	R42	9672	R84	9672	RY4
9672	R44	9672	R85	9672	RY5
9672	R45	9672	R86	9672	RY6
9672	R46	9672	R94		
9672	R51	9672	R95		

Figure 19 (Page 2 of 7). Type and Model values for PROCESSOR Configuration Statement

Type	Model	Type	Model	Type	Model
		Figure 19 (Page 3 of 7). Type and Model values for PROCESSOR Configuration Statement			

Figure 19 (Page 3 of 7). Type and Model values for PROCESSOR Configuration Statement

Type	Model	Type	Model	Type	Model
Amdahl Processors					
0400	GS412	0400	GS422		
0400	GS415	0400	GS425		
0500	GS535	0500	GS555	0500	GS575
0500	GS535E	0500	GS555E	0500	GS575E
0500	GS545	0500	GS565	0500	GS585
0500	GS545E	0500	GS565E	0500	GS585E
0700	GS714	0700	GS7X4	0700	GS765
0700	GS724	0700	GS7Y4	0700	GS775
0700	GS734	0700	GS7Z4	0700	GS782
0700	GS744	0700	GS715	0700	GS785
0700	GS754	0700	GS725	0700	GS795
0700	GS764	0700	GS732	0700	GS7X5
0700	GS774	0700	GS735	0700	GS7Y5
0700	GS784	0700	GS745	0700	GS7Z5
0700	GS794	0700	GS755		
5890	180E	5890	300E	5890	600E
5890	190E	5890	390E		
5890	200E	5890	400E		
5990	250	5990	700	5990	1400
5990	350	5990	790		
5990	500	5990	1100		
5995	250A	5995	3570M	5995	6650M-E
5995	350A	5995	4550M	5995	6670M
5995	500A	5995	4550M-E	5995	7670M
5995	700A	5995	4570M	5995	8650M
5995	790A	5995	4650M	5995	8650M-E
5995	1100A	5995	4650M-E	5995	8670M
5995	1400A	5995	4670M	5995	8671M
5995	2550M-E	5995	5570M	5995	10660M
5995	2570M	5995	5670M	5995	10670M
5995	3550M	5995	6570M	5995	12670M
5995	3550M-E	5995	6650M		
Comparex Processors					
CPX	7/90-1	CPX	7/90-4	CPX	7/90-11
CPX	7/90-2	CPX	7/90-6	CPX	7/90-22
CPX	7/90-3	CPX	7/90-8		

Type	Model	Type	Model	Type	Model
CPX	8/80	CPX	8/90	CPX	8/94
CPX	8/81	CPX	8/90S	CPX	8/95
CPX	8/83	CPX	8/91	CPX	8/96
CPX	8/85	CPX	8/92	CPX	8/98
CPX	8/87	CPX	8/93		
CPX	8/89	CPX	8/93S		
CPX	8/805	CPX	8/830	CPX	8/880
CPX	8/810	CPX	8/840	CPX	8/890
CPX	8/815	CPX	8/850		
CPX	8/820	CPX	8/870		
CPX	9/910	CPX	9/930	CPX	9/950
CPX	9/920	CPX	9/932	CPX	9/960
CPX	9/922	CPX	9/940		
CPX	99/711	CPX	99/822	CPX	99/952
CPX	99/721	CPX	99/831	CPX	99/962
CPX	99/731	CPX	99/832	CPX	99/972
CPX	99/811	CPX	99/941	CPX	99/982
CPX	99/821	CPX	99/942		
CPX	8/805	CPX	8/830	CPX	8/880
CPX	8/810	CPX	8/840	CPX	8/890
CPX	8/815	CPX	8/850		
CPX	8/820	CPX	8/870		
CPX	9/810	CPX	9/830	CPX	9/930
CPX	9/815	CPX	9/840	CPX	9/932
CPX	9/820	CPX	9/910	CPX	9/940
CPX	9/825	CPX	9/920	CPX	9/950
CPX	9/822	CPX	9/922	CPX	9/960
CPX	99/711	CPX	99/822	CPX	99/952
CPX	99/721	CPX	99/831	CPX	99/962
CPX	99/731	CPX	99/832	CPX	99/972
CPX	99/811	CPX	99/941	CPX	99/982
CPX	99/821	CPX	99/942		
CPX	C2000113	CPX	C2000624	CPX	C2000325
CPX	C2000114	CPX	C2000724	CPX	C2000326
CPX	C2000213	CPX	C2000824	CPX	C2000426
CPX	C2000214	CPX	C2000924	CPX	C2000526
CPX	C2000313	CPX	C2000A24	CPX	C2000626
CPX	C2000314	CPX	C2000A25	CPX	C2000726
CPX	C2000324	CPX	C2000126	CPX	C2000826
CPX	C2000424	CPX	C2000225	CPX	C2000926
CPX	C2000524	CPX	C2000226	CPX	C2000A26

Type	Model	Type	Model	Type	Model
		Figure 19 (Page 5 of 7). Type and Model values for PROCESSOR Configuration Statement			
CPX	M2000115	CPX	M2000325	CPX	M2000528
CPX	M2000211	CPX	M2000413	CPX	M2000625
CPX	M2000213	CPX	M2000415	CPX	M2000528
CPX	M2000215	CPX	M2000418	CPX	M2000725
CPX	M2000225	CPX	M2000423	CPX	M2000728
CPX	M2000313	CPX	M2000425	CPX	M2000825
CPX	M2000315	CPX	M2000428	CPX	M2000828
CPX	M2000323	CPX	M2000525		
		Hitachi Processors			
HDS	GX8110	HDS	GX8310	HDS	GX8424
HDS	GX8112	HDS	GX8312	HDS	GX8520
HDS	GX8114	HDS	GX8314	HDS	GX8524
HDS	GX8210	HDS	GX8320	HDS	GX8620
HDS	GX8212	HDS	GX8324	HDS	GX8624
HDS	GX8214	HDS	GX8412	HDS	GX8724
HDS	GX8220	HDS	GX8414	HDS	GX8824
HDS	GX8224	HDS	GX8420		
HDS	GX6110	HDS	GX8210	HDS	GX8412
HDS	GX6115	HDS	GX8212	HDS	GX8414
HDS	GX6210	HDS	GX8214	HDS	GX8420
HDS	GX6215	HDS	GX8220	HDS	GX8424
HDS	GX6225	HDS	GX8224	HDS	GX8520
HDS	GX6325	HDS	GX8310	HDS	GX8524
HDS	GX6425	HDS	GX8312	HDS	GX8620
HDS	GX8110	HDS	GX8314	HDS	GX8624
HDS	GX8112	HDS	GX8320	HDS	GX8724
HDS	GX8114	HDS	GX8324	HDS	GX8824
HDS	PILOT14	HDS	PILOT34	HDS	PILOT75
HDS	PILOT14S	HDS	PILOT35	HDS	PILOT77
HDS	PILOT15	HDS	PILOT35H	HDS	PILOT85
HDS	PILOT15S	HDS	PILOT36	HDS	PILOT87
HDS	PILOT17	HDS	PILOT37	HDS	PILOT95
HDS	PILOT24	HDS	PILOT45	HDS	PILOT97
HDS	PILOT24S	HDS	PILOT47	HDS	PILOTA5
HDS	PILOT25	HDS	PILOT55	HDS	PILOTA7
HDS	PILOT25S	HDS	PILOT57	HDS	PILOTB5
HDS	PILOT26	HDS	PILOT65		
HDS	PILOT27	HDS	PILOT67		

Type	Model	Type	Model	Type	Model
Figure 19 (Page 6 of 7). Type and Model values for PROCESSOR Configuration Statement					
HDS	SKY-11	HDS	SKY-211	HDS	SKY-427
HDS	SKY-21	HDS	SKY-213	HDS	SKY-525
HDS	SKY-22	HDS	SKY-215	HDS	SKY-527
HDS	SKY-31	HDS	SKY-225	HDS	SKY-625
HDS	SKY-32	HDS	SKY-313	HDS	SKY-627
HDS	SKY-41	HDS	SKY-315	HDS	SKY-725
HDS	SKY-42	HDS	SKY-325	HDS	SKY-727
HDS	SKY-52	HDS	SKY-413	HDS	SKY-825
HDS	SKY-62	HDS	SKY-415	HDS	SKY-827
HDS	SKY-72	HDS	SKY-417		
HDS	SKY-82	HDS	SKY-425		
National Advanced Systems Processors					
NAS	EX-10	NAS	EX-40	NAS	EX-95
NAS	EX-11	NAS	EX-42	NAS	EX-100
NAS	EX-20	NAS	EX-44	NAS	EX-210
NAS	EX-22	NAS	EX-50	NAS	EX-220
NAS	EX-25	NAS	EX-60	NAS	EX-310
NAS	EX-27	NAS	EX-65	NAS	EX-310A
NAS	EX-30	NAS	EX-70	NAS	EX-320
NAS	EX-31	NAS	EX-75	NAS	EX-420
NAS	EX-33	NAS	EX-80	NAS	EX-420A
NAS	EX-35	NAS	EX-85	NAS	EX-520
NAS	EX-38	NAS	EX-90	NAS	EX-620
NAS	XL-50	NAS	XL-60M	NAS	XL-80
NAS	XL-50M	NAS	XL-70	NAS	XL-90
NAS	XL-60	NAS	XL-70M	NAS	XL-100
Olivetti Processors					
OLV	PILOT14	OLV	PILOT34	OLV	PILOT75
OLV	PILOT14S	OLV	PILOT35	OLV	PILOT77
OLV	PILOT15	OLV	PILOT35H	OLV	PILOT85
OLV	PILOT15S	OLV	PILOT36	OLV	PILOT87
OLV	PILOT17	OLV	PILOT37	OLV	PILOT95
OLV	PILOT24	OLV	PILOT45	OLV	PILOT97
OLV	PILOT24S	OLV	PILOT47	OLV	PILOTA5
OLV	PILOT25	OLV	PILOT55	OLV	PILOTA7
OLV	PILOT25S	OLV	PILOT57	OLV	PILOTB5
OLV	PILOT26	OLV	PILOT65		
OLV	PILOT27	OLV	PILOT67		

Type	Model	Type	Model	Type	Model
OLV	SKY-11	OLV	SKY-211	OLV	SKY-427
OLV	SKY-21	OLV	SKY-213	OLV	SKY-525
OLV	SKY-22	OLV	SKY-215	OLV	SKY-527
OLV	SKY-31	OLV	SKY-225	OLV	SKY-625
OLV	SKY-32	OLV	SKY-313	OLV	SKY-627
OLV	SKY-41	OLV	SKY-315	OLV	SKY-725
OLV	SKY-42	OLV	SKY-325	OLV	SKY-727
OLV	SKY-52	OLV	SKY-413	OLV	SKY-825
OLV	SKY-62	OLV	SKY-415	OLV	SKY-827
OLV	SKY-72	OLV	SKY-417		
OLV	SKY-82	OLV	SKY-425		

Figure 19 (Page 7 of 7). Type and Model values for PROCESSOR Configuration Statement

**** Measured Usage Joint Study Technical Document end ****