Compuware | The Mainframe Software Partner For The Next 50 Years

ThruPut Manager
Drive Booking Services (DBS)
System Programming Guide

Release 18.02

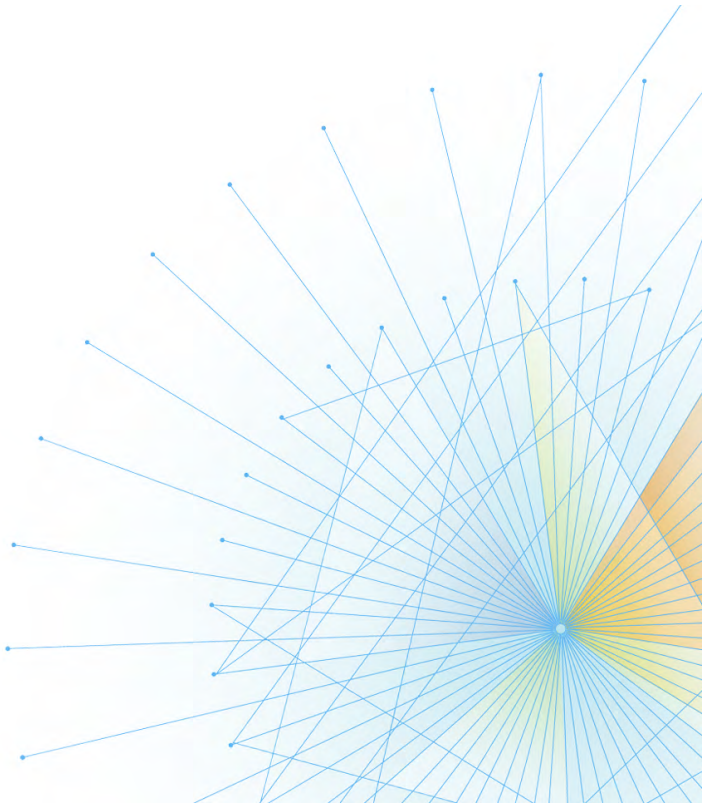

Please direct questions about ThruPut Manager
or comments on this document to:

ThruPut Manager **Customer Support**

https://go.compuware.com/

Doc. APR2019
February 22, 2019

Introduction

Summary of Changes

V1802-7118 (*April 2019*)	• No changes
V1802-7117 (*January 2019*)	• No changes
V1802-7116 (*October 2018*)	• No changes
V1802-7115 (*July 2018*)	• No changes
V1802-7114 (*January 2018*)	• No changes
V1802-7113 (*January 2018*)	• No changes
V1802-7112 (*October 2017*)	• No changes
V1802-7110 (*July 2017*)	• Rebranding of MVS Solutions to Compuware. This includes update of cover style, copyright, and changing version release to 18.02.
V7R1-7109 *(April 2017)*	• No changes
V7R1-7107 *(May 2016)*	• No changes
V7R1-7106 *(November 2015)*	• No changes
V7R1-7104 *(July 2015)*	• No changes
V7R1-7101 *(July 2014)*	• This is a base manual for ThruPut Manager Version 7 Release 1.0.

About This Manual

This manual provides information for systems programmers involved in planning and customizing the Drive Booking Services (DBS).

Contents

Introduction . . . 3
Summary of Changes . . . 3
About This Manual . . . 3

Chapter 1 The DBS Function . . . 9
The Drive Booking Services Component . . . 9
Functions Provided With DBS . . . 9
DAL and JAL for DBS . . . 9
JECL for DBS . . . 9
Operator Commands for DBS . . . 10
TMISPF Dialogs and Displays . . . 10
Expanding and Collapsing Branches . . . 10
The Action Bar . . . 10
DBS Objectives . . . 10
Overview of DBS . . . 11

Chapter 2 DBS Concepts . . . 13
The DBS View of Tape Subsystems . . . 13
Drive Pools . . . 14
Work Groups . . . 14
The DBS View of Your Installation . . . 15
Counts Versus Device Numbers . . . 15
The Configuration and Policies . . . 15
DBS Priorities . . . 16
DBS and Started Tasks/Jobs . . . 16
Batch Jobs and Dynamic Allocation . . . 17
New Descriptors in JAL . . . 17

Chapter 3 How DBS Works . . . 19
DBS Job Flow Management . . . 19
The Job Selection Process . . . 19
Execution Time Actions . . . 19
DBS Job Status . . . 20
The DBS System Affinity Mask . . . 21
Other Considerations . . . 22
A Schematic Description of DBS Mask Management . . . 23
Managing the DBS System Affinity Assignment . . . 25
OPENing AND CLOSing Pools . . . 25
DBS Unavailable Devices . . . 26

Chapter 4 Planning for DBS . . . 27
Introduction . . . 27

The Automation File . . . 27
Configuration Definition . . . 28
Policy Planning . . . 28
Work Group Planning . . . 29
What is a Work Group? . . . 30
Work Groups Versus Drive Pools . . . 30
Apportioning Drives Using Work Groups . . . 31
Implementing Work Groups . . . 31
The Definition Process: What Should I Have? . . . 33
DBS and Job Limiting (JLS) . . . 36
The Old Role of JLS . . . 36
The New Role of JLS . . . 36
An Example of Using JLS to Extend DBS . . . 37
Relating DBS Descriptors to $UNIT Descriptors . . . 38

Chapter 5 Implementing DBS . . . 39
System Requirements . . . 39
DBS Facilities . . . 39
Installation and Implementation . . . 39
DBS Configuration Definition Considerations . . . 39
Invoking the DBS Configuration Dialog . . . 40
JESplex Configuration . . . 41
Defining the First Configuration . . . 41
The **BASE** Policy . . . 41
DBS Work Groups . . . 42
SMF Monitoring of DBS Activity . . . 42
Activating Data Collection for DBS . . . 42
Identifying DBS Records . . . 42

Chapter 6 Administering DBS . . . 43
Introduction to DBS Administration . . . 43
Managing the Configuration . . . 43
Configuration Definition . . . 43
Vendor Selection . . . 44
Selecting a Mode and Naming a Library . . . 44
Selecting Device Types . . . 46
Specifying Device Numbers . . . 46
Activating a Configuration . . . 47
Creating, Editing, and Replacing the NEXT Configuration . . . 47
Backing Out of a Configuration . . . 48
Managing Policies . . . 48
Creating Specialized Policies . . . 48
Setting Device Counts . . . 49
WTO Messages . . . 50
Activating a Policy . . . 50
PASSIVE Mode . . . 50

Configuring Work Groups 50
Renaming Work Groups 51
Setting the Default Work Group 51
Turning Overbooking Off 51
Setting Work Group Counts 51
Using Installation Pools 51
Using JAL and JECL with DBS 52
Assigning Work Groups 52
Setting DBS Priorities 52
Implied Priority 52
Assigning Installation Pools 52
Unavailable Unit Handling 53
JECL for STCs 53
Reserving Devices With //*+DBS RESERVE 53
Setting Work Groups and Priorities With //*+DBS SET 53
Accounting for Dynamic Allocation 53
DBS Monitor and Displays 54
DBS Display Services 54
DBS Monitor Services 54
DBS Configuration/Policy Report 54

Appendix A Supported Devices 59

Chapter 1
The DBS Function

This chapter describes the functions provided by the Component Drive Booking Services.

The Drive Booking Services Component

ThruPut Manager offers a Component that is installed on top of the ThruPut Manager Base Product. The Component is named “Drive Booking Services”, or DBS for short. It provides two related but distinct services:

- The management of jobs that require tape drives, real or virtual. The drives that are managed represent all the devices that are available to batch processing in a given JESplex.
- The subdividing of the available tape drives into Work Groups to satisfy the requirements of different types of work, for example production versus non-production.

This chapter introduces the different functions provided with this package.

Functions Provided With DBS

This Component includes the following additional functions to support DBS:

- DAL/JAL extensions.
- JECL Service Extensions.
- Operator commands.
- TMISPF displays and dialogs.
- Initialization statement changes.

DAL and JAL for DBS

DBS supports the JAL Action statement DBS SET, allowing you to set the Work Group and Priority for a job. The DBS HOLD statement lets you specify that a job should be held rather than failed if it requests more device than are available in the ACTIVE Configuration.

If Installation Pools are used, the DAL Action statement DBS ASSIGN lets you indicate that a device belongs to a particular Installation Pool.

JAL Descriptors allow you to display in a message the DBS settings for a job.

For details and syntax descriptions, refer to the *DAL Reference Guide* and the *JAL Reference Guide*.

JECL for DBS

ThruPut Manager’s Job Analyzer cannot determine the correct counts for tape devices needed by jobs that use dynamic allocation for tapes. Also, the Job Analyzer does not examine STCs. You can, however, still have DBS manage the drives for STCs and jobs that use dynamic allocation. Simply add //*+DBS RESERVE JECL statements describing the requirements for each affected step.

For example:

```
//*+DBS RESERVE IBM->AUTOMATED->ATLDS1->3480S->3480=2
```

This JECL statement allows DBS to maintain accurate counts for tape device usage.

You can also use JECL to specify the Work Group and/or DBS Priority for a job or STC by using the //*+DBS SET JECL statement:

```
//*+DBS SET WORKGROUP=SPECIAL,PRIORITY=HIGH
```

For details and syntax descriptions of these statements, refer to the *System Base Product: Programming Guide*.

Operator Commands for DBS

Operator commands allow you to:

- Verify a DBS Configuration (DBS VERIFY).
- Re-evaluate the ACTIVE Configuration if your I/O EDT changes (DBS REEVALUATE).
- Activate Policies (DBS ACTIVATE).
- Alter the DBS priority for a job (DBS ALTER).
- Display the current status of DBS (DBS DISPLAY).
- Change the DBS status of a DBS-managed device (DBS SET).

For details and syntax descriptions, refer to the *Command Reference Guide*.

TMISPF Dialogs and Displays

DBS uses dialogs in TMISPF to create and manage the Configuration and Policies. (For information about TMISPF, see the *System Programming Guide: Base Product*.)

TMISPF also displays information about the jobs and Drive Pools managed by DBS. Before using them, you should be familiar with some conventions.

Expanding and Collapsing Branches

Some windows use a display in the form of a tree structure. Lower levels can be expanded and collapsed. A plus sign (+) indicates that the branch is collapsed, and a minus sign (-) indicates that the branch is expanded. You expand or collapse branches by placing the cursor within the Pool name and pressing **Enter**. You will see this tree structure used in several DBS management dialog panels.

The Action Bar

When using the DBS dialogs, do not overlook the Action Bar:

- The **Edit** menu allows you to change such things as the Work Group settings and information.
- The **GoTo** menu is a useful way to go directly to the panel in which you are interested.
- Help is self-explanatory.

Note that not all items are available in every panel.

> Important information regarding the UNSYNC Indicator: When the Configuration you are editing does not match the JESplex on which you are running, the status is shown in the upper right of the screen as UNSYNC.

DBS Objectives

DBS has been designed to provide a completely automated solution to the problem of allocating a serially reusable resource: tape drives.

The only demand that DBS makes on your installation before it can function is that you must define the tape subsystem environment to DBS.

This definition is a simple process that is assisted by a comprehensive ISPF dialog. If your environment is simple, then the definition process is extremely simple. If your environment is complex (many different devices and lots of asymmetries) then the process is a bit more elaborate.

The point here is that the complexity of the definition process is related to the complexity of your environment. Fortunately, the ISPF dialog acts as a "complexity reducer", dampening the magnification effect that complexity creates.

Although the algorithms and processes used to manage your installation's tape drives are complex, DBS is straightforward to set up, and once implemented requires little intervention.

Overview of DBS

The first function of DBS is to manage the totality of drives available to the JESplex.

The second function, which is optional, allows the logical separation of drives to provide resources to different work groups. Of course, subdividing must take place within the boundaries of the available drives.

To manage the drives, DBS uses these resources:

- *Drive Pools* manage the physical devices (device numbers).
- *Work Groups* manage the logical segmentation of workload.

Externally, DBS requires a user-defined Configuration and Policies to determine how your installation wants to manage tape drives. These are defined with ISPF dialogs. DBS also uses ISPF to provide monitoring and display facilities.

The Configuration, Policies, Drive Pools, Work Groups, and how they fit together are described in more detail in the following chapters.

Chapter 2
DBS Concepts

This chapter explains certain concepts of DBS that are important to the planners and implementers for ThruPut Manager. Some of these concepts are explained in the *DBS Concepts and Facilities* manual, but for convenience we repeat them here.

The DBS View of Tape Subsystems

Before we describe Drive Pools, we first need to explain the structure of the different types of drives, modes of operation of each drive, and the different vendors as seen by DBS. There is nothing particularly unique about DBS's view but since several choices were available, you should understand the DBS approach.

The DBS view of tape drives is consistent with the way the software support for this type of hardware is generated and referred to. The lowest level is what we call a Drive Pool. The aggregation of Drive Pools is done in four hierarchical levels:

- Level 1 (highest level) represents the IBM device types, for example 3590. All devices *generated* with the generic name 3590-1 (regardless of what physical hardware they truly represent) are aggregated at Level 1 for the 3590 device type.
- Level 2 represents the devices associated with a particular vendor. For actual physical drives we have two distinct vendors: IBM and StorageTek. For Virtual Tapes we have two additional vendors with a software solution: CA-Vtape and COPYCROSS.
- Level 3 represents the mode each vendor uses to manage the drives: Manual, Automated, and Virtual.
- Level 4 represents the type of drive. It has a level and an associated sublevel. The *level* shows how the system recognizes the device, for example 3480. (This is the same device that is aggregated at Level 1 and is found in the EDT.)
  - The *sublevel* represents the actual type of drive. For example, for an IBM 3590-1 the sub-device could be a 3590H or a 3590E. For StorageTek, a 3590-1 could be a T9940A35. Standard system allocation is unaware of the sublevel.

To summarize:

- Level 1: Total number of devices for each generic device type, regardless of vendor, mode of operation, or type of drive.
- Level 2: The Vendor (IBM, StorageTek, COPYCROSS, CA-Vtape).
- Level 3: Mode of Operation (Manual, Automated, Virtual).
- Level 4: Generic Device type (e. g. 3590-1).
  - Sublevel: Actual device (e. g. 3590H).

When defining drives to DBS, you have to be concerned with only the lowest possible level, that is, the Sublevel for Level 4. This represents the device type that the person creating this type of definition is familiar with. DBS then calculates the aggregated values for all other levels.

These levels dictate the sequence in which you define your devices to DBS:

1. You select the vendor.
2. You select the mode of operation.
3. You define your libraries by naming them and selecting the generic device types.

4. You associate specific device numbers with the device types.

Once you grasp this structure, it is relatively easy to understand DBS Drive Pools.

Drive Pools

A Drive Pool represents the lowest possible group that can satisfy a particular allocation. Any device in the group is interchangeable for the purpose of non-directed allocation (directed allocation takes place when a request is made for an specific device number, for example '0380'). To illustrate the point we can use IBM's Virtual Tape Systems (VTS).

If an installation has only one VTS subsystem then the answer is obvious. All the device numbers associated with the VTS represent a unique DBS Pool.

With two VTS subsystems, even though the devices are homogeneous (all 3490E) the device numbers are not interchangeable across VTS subsystems. This is simply the result of a set of device addresses not having access to a set of volumes. So, at least for specific volume requests, we have two distinct DBS Drive Pools.

In the case of a non-specific request, depending on the installation's SMS definitions, prior to allocation it might not be possible to determine what DBS Drive Pool will be used to satisfy the allocation. Here is where the next DBS Level plays a role. DBS knows that the allocation will be satisfied by a VTS, so it uses the values of the "aggregated" Virtual Drive Pools until the allocation takes place. Once the allocation has occurred, the actual DBS Drive Pool is known. At that time (immediately after allocation) DBS does the necessary adjustment to reflect the actual DBS Drive Pool.

Work Groups

DBS Work Groups allow the logical segmentation of your workload so that you can control how different types of jobs compete for the devices in a Drive Pool. In simple terms this could be your Production versus Non-Production workloads. This service is optional.

For each Work Group, you assign minimum and maximum values for every Drive Pool. This is organized in a two-level structure, represented in the following DBS work group structure diagram.

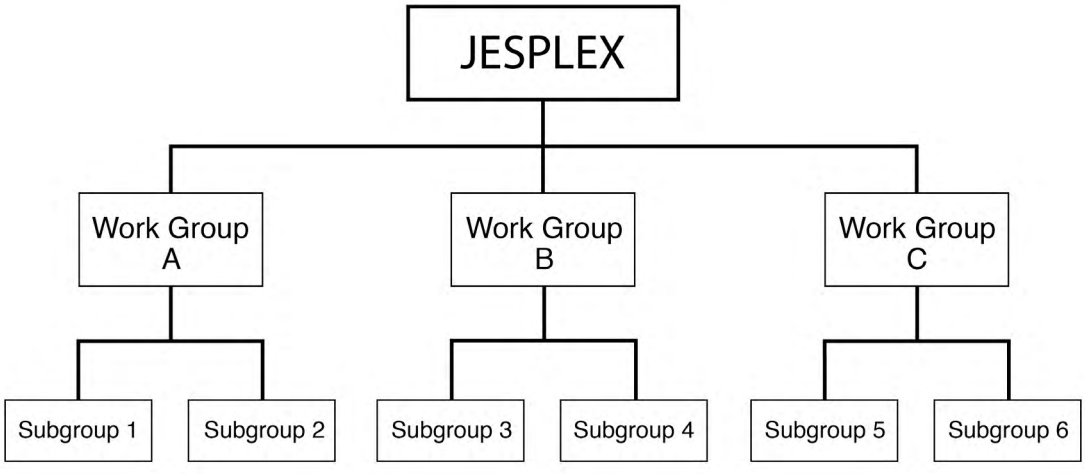

What this structure means is that you can first segment your workload into three possible groups. For example: Production, Non-Production, Development. Each of these three groups can then be subdivided into two groups.

What is the purpose of these divisions? In essence it is to "apportion" the existing drives in a manner that recognizes the type of work in question. For simplification, let's assume the following:

- There are only two possible groups: Production, Non-production.
- There are 64 available drives.

If you do not define Work Groups then Production and Non-Production will have up to 64 drives at their disposal and they will contend for them. With Work Groups you could decide that Production can allocate any of the 64 drives but Non-production is only allowed 32 drives. In this way you have created an environment where Production has *32 drives for their sole use* and another 32 to contend with Non-Production jobs. Non-Production does not have dedicated drives but can contend for up to 32 drives with Production.

Since there is a set of Work Group values for every Drive Pool, you can view the Drive Pool as the description of what is available and the Work Groups as "who gets what."

The DBS View of Your Installation

DBS manages the tape drive counts for a particular JESplex. The reason for this scope is that batch is managed at the JESplex level. Once the JESplex is identified, the participating systems are then identified. This is for the purpose of verification.

The tape subsystems are defined first to the JESplex. Here, all the possible universe of device numbers that can be seen by any of the participating systems is defined.

By design, the participating systems inherit all the tape drives that are defined to the JESplex. This assumes that you have a symmetric environment:

- If this is the case, then no device specification is necessary at the system level.
- If the environment is not symmetric, then the drives that are not visible to a particular system are removed from that system's list.

DBS supports drives that have one device number on one system but a different device number on another system. It does not support environments that have different devices with the same device number in different systems. Please note:

> DBS does not perform tape switching or sharing.

The installation is expected to have a solution to this problem. Either IBM's Tape Switching facility or CA's Tape Sharing facility (or any other similar facility) is expected to be in place in environments that share tape drives across multiple systems.

Counts Versus Device Numbers

It should be noted that *all this management is done with "counts."* DBS does not direct allocation to particular device numbers. The device numbers are provided to DBS for the purpose of verification after allocation has done its work. This approach greatly simplifies the facility and eliminates any possible conflict between system allocation's opinion and that of DBS. *Allocation is always right*, so when DBS finds itself in disagreement with allocation, it discreetly changes its mind.

> It is important to understand that the device number list can contain more devices than the Drive Pool count. This means that if the system allocates any of the device numbers in the list, the in-use count for its associated Drive Pool is updated.

The Configuration and Policies

In order to manage your installation's tape drives, DBS uses a Configuration and one or more Policies:

- A *configuration* provides DBS with a description of the physical and virtual drives to be managed.

- A *policy* describes to DBS how to manage the resources defined in the Configuration.

DBS needs to know what is available to the JESplex *under its management*. DBS cannot assume that because a device is "visible" to a system it is to be available for batch processing. So the simplest approach is for the installation to inform DBS of what is under its control. This is called a DBS Configuration. It includes all the maximum inventory available to the JESplex. In fact, you can (and should) include future devices in the Configuration.

Within a given Configuration, an installation can create Policies. In a Policy, the installation describes how the tape drives (or a subset) are to be deployed.

You change a Configuration only when your tape subsystems are being modified, for example, when a bank of drives is about to be removed. Think of the Configuration as a static framework in which Policies operate.

Policies are designed to be as fluid as the installation needs them to be. When you define a Configuration, you automatically define your initial Policy, which is called "**BASE**". This Policy assumes that all systems in the JESplex share all defined drives unless you edit it.

You can have only one Configuration and one Policy active at any given time in the JESplex. Activating a new Configuration is a more restricted action than activating a Policy. The removal of devices can impact jobs that are already in the system, so DBS provides facilities to ensure effective transitions that will not result in jobs being failed. On the other hand, Policies can be activated and deactivated at any time.

DBS Priorities

For jobs that are not managed by SLM, DBS allows you to assign a priority through JAL or JECL:

- LOW
- MEDIUM, which is the default.
- HIGH

When contention occurs, the DBS Priority determines which step gets the drive. See Setting DBS Priorities.

DBS and Started Tasks/Jobs

Started tasks that allocate drives (defined to DBS) can run in one of two modes:

- As DBS participants
- As poachers

This is not an all-or-nothing situation. You can run some in one mode and others in the other mode. If you do not take any actions, they are treated as poachers.

DBS regularly reviews the status of the drives it is managing. When it finds out that some of the devices it is managing are allocated to a non-registered task it treats that task as a poacher. It collects as much information as possible and records it in the poachers section.

The count adjustments for poachers are done differently from DBS registered users. When a drive is confirmed, after allocation, to have been allocated to a registered user the following takes place:

- The available count is reduced.
- The in-use count is increased.
- The total number of drives is not altered (available + in-use).
- The device number is flagged as in-use.

When a drive is found to have been poached, the following takes place:

- The count of device numbers in the defined list is compared to the total count.

- If there are more drives in the device number list than in the count (the reverse cannot happen), then:
  - The particular device number is flagged as in-use by a poacher.
  - The device number will be associated with the poaching task until DBS detects that it is no longer allocated.
  - No other adjustments are made.
- If the count of the device number list matches the total count, then:
  - The total count is reduced. It is treated as if DBS has lost a drive.
  - The available count, by definition, is equally reduced.
  - The device number is flagged as in-use by a poacher.
  - DBS will be short of one drive until the poacher gives up the drive.

The installation might prefer to run the started tasks as a DBS participant. A JECL statement is available for that purpose. The type and number of devices can be specified for each step that dynamically allocates drives. For each specified Pool, two values can be coded:

- One value represents the static allocation needs.
- The other value represents the dynamic allocation needs.

The second value results in DBS reserving the equivalent number of drives for the duration of the step. For started tasks that cannot handle dynamic allocation failures, the RESERVE mechanism is a good solution.

Batch Jobs and Dynamic Allocation

You can let DBS know that a particular job/step allocates drives dynamically. As indicated earlier, DAL directives and a JECL statement are available for that purpose.

The JECL statement has the same format as the one to be used with started tasks. The only difference is that of the two values that can be coded as the count of drives needed, only the first one has meaning. The second value represents static allocation needs for a started task. It is not needed for batch jobs because the static count is calculated at analysis time. If you code a value in the second field for batch jobs it is ignored.

New Descriptors in JAL

A new set of job descriptors have been added to JAL. They more closely parallel the DBS Pools than the traditional UNIT counts. There is a section later in this publication, DBS and Job Limiting (JLS) that outlines the difference between the old UNIT descriptors and the new DBS Pool Descriptors.

Chapter 3
How DBS Works

This chapter reveals some of the internal algorithms and methods used by DBS.

DBS Job Flow Management

DBS adds intelligence to the job selection and step initiation/termination process; however, it does not decide which job is "next." That choice is left to either the traditional JES2 initiators or WLM initiators. ThruPut Manager might assist these initiators, but the fundamental job selection process is not altered.

When DBS is running under the control of SLM, it is SLM's decision which job is "next." What DBS does is to determine whether or not the job should be allowed to proceed to initiation based on the availability of tape drives. DBS either accepts the job or rejects it. If the job is rejected, selecting "next" is again done by the facilities mentioned above, not by DBS.

A job must be analyzed in order to be under DBS management. That is where the job's requirements for tape drives are determined. As a result of the analysis process, a Drive Pool mask and a Work Group mask are created and associated with the job. DBS is aware of all the defined Drive Pools and Work Group Pools and has assigned a bit for each one to be used when constructing the masks. The job masks are identical in structure to the JESplex/Member masks. In addition to the masks, all the "counts" and the high watermark for the job's requirements are also calculated and associated with the job.

The Job Selection Process

The DBS job selection process is not based on "counts." The concept is different. Each defined Pool is either OPENED for additional jobs or it is CLOSED. If any of the Pools required by a job are CLOSED, then the job is rejected. This decision is a very efficient process using the masks. OPENing and CLOSing Pools is a separate process that involves several considerations, such as the sum total of all the high watermarks of jobs in execution, overbooking factors, steps waiting for DBS "permission" to enter allocation, and historical trends.

DBS maintains a JESplex mask and a mask for each participating JES2 member. The member masks, as you would expect, are each a subset of the JESplex mask. The reason that a mask is needed for individual members is to allow for asymmetries. This can occur at the Configuration level or at the Policy level. Configuration asymmetries tend to be "hard" ones. Policy asymmetries are usually logical ones.

Execution Time Actions

Once a job is selected, DBS tracks the job step by step until it terminates. DBS is aware of all the requirements for each step. It also keeps track of the initial high watermark and subsequently reduced high watermark. The drive requirements for a step were calculated during the simulation/analysis process. DBS refers to these values as "best effort" determination. The term is intended to highlight the possibility that actual allocation might produce different results. *For obvious reasons, allocation is always right.* So when discrepancies are encountered, the "best effort" value is changed to reflect the actual value.

The actual value is determined immediately after allocation. All steps are scrutinized, even when the DBS "best effort" did not find any tape drive requirement. The actual results are compared to the "best effort" values and, if necessary, adjustments are made.

If drive allocation occurs for a job for which DBS did not find any requirements, DBS is informed. Since the job was not expected to need any drives, it was not registered with DBS at job initiation. When actual drive requirements are found the job is then registered with DBS, so any job that was analyzed is registered with DBS when executing regardless of the best effort determination.

For steps for which drive allocation is expected, a call is made to DBS to determine if the requirements can be satisfied before it is allowed to proceed to actual allocation. This approach prevents steps from entering allocation recovery; however, the more important reason for jobs to wait prior to actual allocation is to allow DBS priority management, not only for a system but also across the JESplex. Once a step enters allocation recovery, there is only a limited amount of control that can be exercised. The HOLD or NOHOLD facility is system oriented and in general is first-in-first-out.

The step call to DBS may be viewed as a request for "permission" to enter allocation. If DBS determines that the requirements can be satisfied, it then:

- Adjusts the counts to reflect the needs as determined by "best effort."
- Posts the caller so it can proceed to allocation.

If DBS determines that the requirement cannot be satisfied, it then:

- Records the "best effort" requirements for a job/step. It also records the DBS priority.
- Denies the step permission to enter allocation.

When drives become available (normally at step termination of other jobs), DBS asks:

- Is any step in the JESplex waiting for "permission"?
- Can the newly available drives satisfy the requirements of the highest priority step?

If the answer is YES then it will post the appropriate step or steps. This is the mechanism to honor DBS priority. Steps with equal priority are handled first-come-first- served.

After static allocation is done, a call is made again to inform DBS of the actual drive counts that have been allocated. DBS then does the necessary adjustments by comparing "best effort" to actual allocations. Certain type of discrepancies are logged. This allows a review of what might be causing the differences so that "best effort" calculations can be improved.

It should be noted that one of the triggers for DBS to CLOSE a Pool or Pools is the denial of "permission" for any step to enter allocation. The Pool or Pools causing the delay are immediately CLOSED.

DBS is notified when a step has completed execution. It then:

- Returns the drives to available status.
- Determines if the high watermark for the corresponding job has to be adjusted.
- Determines if any step is waiting for the type and number of drives that are now available. If that is the case, it posts the step or steps.

DBS is also notified when a job terminates. At this point the only action required is to de-register the job.

DBS Job Status

A DBS managed job can be in one of the following states:

- In the JES2 execution queue waiting to be selected. Its turn has not arrived yet.
- In the JES2 execution queue waiting for a DBS Pool or Pools to be OPENED:
  - If the Pool is CLOSED at the JESplex level it is the equivalent of the job being in HOLD status.

  - If the Pool is CLOSED for a particular system, it is the equivalent to the job having affinity to the systems that have the needed Pools OPENED.

- In the JES2 execution queue with no DBS system affinity to any system because the job exceeds the JESplex drive resources. This situation is Policy dependent. It is a "conceptual" HOLD since the job is not in a hard HOLD. It is similar to the status of a job when no scheduling environment is active that can satisfy the job requirements.

- The job could be executing a step with no drive requirements. In this state there no direct relationship between the step and availability of drives. DBS knows about future requirements for the job and that information is used by the overallocation algorithm, but it has no direct bearing on the values for Pools.

- The job could be executing a step that requires tape drives. In that case there is a direct relationship between the step and the count values of the Pools the step needs. Any drives the step is using were moved from available status to in-use status and will stay that way until the step terminates.

- The job could be waiting in a step for DBS "permission" to enter allocation. DBS is aware of the step requirements and its priority. When resolving the situation it does not do best fit. The highest priority/earliest arrival step is served first. No drives are given from that Pool to any step request in the JESplex until the top priority waiting step can be satisfied. Steps can jump ahead of other steps if they have higher priority; however, this will not cause the lower priority steps to wait "forever." Since the Pool or Pools needed are CLOSED there is only a finite number of steps running in the JESplex. Before any other job is allowed to initiate, they will all be served as other steps terminate and drives are returned to the available pool.

- Under unusual circumstances a job could be in allocation recovery. This could occur under the following conditions:

  - DBS only controls jobs that are DBS-managed. It does not serialize other tasks. It is possible for DBS to grant "permission" to a step to enter allocation, but between the time permission is granted and allocation begins for the step a poacher acts. There is a window of opportunity for a poacher to "steal the drives." The more poachers you have running in high priority service classes, the more likely this is to occur.

  - The other situation is when "best effort" under-counts the number of drives needed. There might not have been any additional available drives when permission was granted, but if the actual requirement had been known the step would have been placed in a wait status.

Once the job terminates DBS does not keep any historical information about that particular job. It accumulates historical trends to calculate overallocation factors, but no job-specific details. If you need downstream data about the jobs managed by DBS or about tape drive utilization, you can request DBS to generate SMF records.

The DBS System Affinity Mask

One more thing... (Does it ever end?) There is an additional consideration in the job selection process: if you have asymmetric systems (as a result of a Policy or Configuration) the following situation could occur:

- A particular system has a significantly reduced drive count. (If it is zero then there is no problem because the Pool will always be CLOSED in that system). As a result there might be jobs that exceed the available resources for that system.
- These jobs cannot be run in that system (or systems).
- When there are drives available (it could be as few as 1) the Pool will be OPENED.
- Unless something additional is done to prevent the jobs that exceed requirements from being selected, the "is this Pool OPENED?" job selection logic will allow them to continue to job initiation. This is not a very good idea.

To handle the situation described above DBS, in addition to the Drive Pool masks, constructs a DBS System Affinity mask. Jobs that exceed the requirements of a particular system will not have DBS affinity to that system. The first thing that DBS does during the selection process is to test the affinity.

In addition to asymmetric systems you can have a situation that a Policy is loaded with a significantly reduced number of drives for the JESplex. This could be the result of hardware out of service, or rearrangement of availability of drives across JESplexes at different times of the day. Any job that exceeds the resources made available by the Policy for the JESplex will have all its DBS System Affinity removed so it will not be selected by any system.

Note that if the Configuration for the JESplex cannot satisfy the requirements for a particular job, this is a "will never be able to run" situation regardless of the Policy, since a Policy cannot exceed the resources defined in a Configuration. This situation usually calls for an informative message and a JCL error, but if desired, you can use the JAL statement DBS HOLD to request that the job be placed in HOLD in the MHS_TM category.

The DBS System Affinity mask should not be confused with the OPEN/CLOSE mechanism used for job selection:

- The DBS System Affinity mask addresses the problem of jobs exceeding existing resources as described by the Policy (either at the system level or JESplex level).
- OPEN/CLOSE manages the normal job selection to maximize drive usage without overallocation problems.

Other Considerations

DBS is also called during step execution in two situations:

- A dynamic allocation has taken place that required tape drives.
- A FREE=CLOSE has caused a drive to be unallocated.

Normally, the job analysis process cannot detect dynamic allocation requirements. If the installation takes no action to let DBS know about the dynamic allocation of drives by a particular job/step, the following takes place:

- Dynamic allocation will enter system allocation without DBS having any knowledge:
- If it can be satisfied, DBS is invoked after allocation has taken place. As a result of this call the in-use counts for the job will be updated and the availability counts will be reduced. As with step allocation, the actual job might have to be registered if there were not any apparent drive requirements at the time of analysis.
- If it cannot be satisfied, the caller receives (in most cases) an "unable to allocate" return code from dynamic allocation. It is then up to the caller to decide what to do. In a number of cases that can result in a step failure. In other cases, the executing code might have a wait-and-try-again mechanism based on some reasonable time interval. The behavior of a particular step is situational and normally only known to the installation. No generalizations can be made, other than "if drives are always available things are always OK."

For situations where the unavailability of a tape transport might result in a step failure, DBS provides a facility to prevent that occurrence. As you might expect the need has to be externalized. A job/step can request, either in DAL or JECL, that a number of drives of a particular type be RESERVED. This is in addition to the ones needed and externalized in JCL. The step is not given "permission" to proceed to allocation until the sum total of the "best effort" count and the number requested to be RESERVED are available. The RESERVED count will continue to be in effect until the step terminates.

If and when dynamic allocation occurs, and the actual drives are of a different type from the ones that were RESERVED, they are treated as additional requests and added to the total.

The case of FREE=CLOSE is a simple one. The appropriate number of drives are moved from in-use to available. The particular step has it in-use count reduced.

A Schematic Description of DBS Mask Management

This section is intended for readers who are more oriented towards schematic descriptions rather than long explanations. It reiterates what the previous sections have described. It focuses on the role played by the DBS System Affinity masks during job selection.

For the purpose of illustration, let's assume the following:

- A JESplex with two participant LPARs (member A and member B).
- A DBS Configuration with 6 Pools defined.
- To simplify the description, the size of the Drive Pool mask allows for up to 8 Pools. The actual mask allows for up to 128 Pools.
- Similarly, the JESplex mask is restricted to 8 systems. The actual mask has room for 32 systems.

When a Policy is activated in a JESplex with no DBS managed jobs running all the Pools are OPENed since there is no usage. DBS will create the masks as follows:

JESplex	1111	1100

MEMBER A	1111	1100
MEMBER B	1111	1100

This represents the 6 Drive Pools that are defined in this installation. All the bits are "ON" so all the Pools are OPENed.

The equivalent Work Group masks are more complex. There are a total of 18 masks. There are 6 Work Groups so each Work Group needs a mask set. 3 sets are needed. One for the JESplexs, one for member A, and one for member B.

WORK GROUP 1	1111	1100
WORK GROUP 2	1111	1100
WORK GROUP 3	1111	1100
WORK GROUP 4	1111	1100
WORK GROUP 5	1111	1100
WORK GROUP 6	1111	1100

If the systems are symmetric, as in the case of the example Policy, the JESplex masks and the individual system masks will be the same. Only when there are system asymmetries would the JESplex masks and individual system masks be different.

When a job arrives—let's call it JOBA—it is analyzed. If it needs tape drives the DBS System Affinity mask, the Device Pool mask, and an index to the Work Group masks (which one to use of the possible 6) are created. Let's say that JOBA needs drives from only one Pool, the second one. The job is assigned to the 4th Work Group. Normally, the installation names the Work Groups. Internally, they are assigned an index value from 00 to 05.

The DBS masks that are constructed for the job are as follows:

DBS System Affinity mask	1111	1100
Drive Pool Mask	0100	0000
WORK GROUP Index (value 03 for the fourth WORK GROUP)	0000	0011

- The DBS job affinity mask indicates that under the active Policy, the job can run on any system. It does not exceed the available resources.

- The Drive Pool Mask indicates that it needs drives from Pool 2.
- For the Work Group, there is no need for a Work Group Pool mask since the requirements are identical to the Drive Pool mask. What is needed is a pointer to the Work Group mask to be used for job selection. For JOBA the value is 3 to represent the 4th Work Group. (Only in the computer world does counting start at zero!)

As jobs arrive and are analyzed, the DBS masks are constructed (among other functions) and then the jobs are placed in the execution queue.

At job selection time, DBS is asked to verify if the selected job should be allowed to proceed to initiation. (Other functions are also called for that purpose.) Let's suppose that the job that has been selected is the one that was initially described, that is, JOBA. This occurs in member B (the second JESplex system). Let's further assume that all Pools are OPENED. DBS will do the following:

- Determine if the job can run in this system. Since the DBS System Affinity mask is like this (shown before):

DBS System Affinity mask	1111	1111

  JOBA can be selected in any system, so the first requirement is satisfied.

- The next verification is related to Work Groups. Here we have a set of 6 masks for the JESplex and a set of 6 masks for member B. From the Work Group index, the correct mask for Work Groups is the 4th one. This is true for the JESplex and for member B. DBS then:
  - Extracts the two masks.
  - The masks are ANDed. The bits must be "ON" in both masks indicating that the Pool is OPENED at the JESplex level and for member B.
  - The JOBA Drive Pool mask is used for the determination of Work Groups and Drive Pools. In this case the Pool masks will look like this:

Job Mask	0100	0000
JESplex and member B mask (ANDed)	11111	1100

  - The Pool needed by JOBA is OPENED for its Work Group, so DBS proceeds to the next and final verification.
- The Drive Pool verification does not require any index since there is only one mask for the JESplex and one mask for member B. The process is identical to the one illustrated above. JOBA, from a DBS point of view, can be allowed to proceed to initiation.

  For JOBA to be affected the 2nd Drive Pool or the 4th Work Group Pool must be CLOSED. Any other activity for other Pools or Work Groups does not affect the eligibility for JOBA's selection.

  To show the role of the DBS System Affinity mask, let's assume that a new Policy has been activated. This Policy, for whatever reason, drastically reduces the number of drives available to member A. At Policy activation time DBS reviews the jobs in the execution queue to see if any one of them exceeds the reduced number of drives available to member A. Any job found to exceed resources will have the member A (system 1) affinity removed, so the mask will look like this:.

Job Mask	0111	1111

None of these jobs will pass the first step in the DBS verification process in member A, so they will not be selected.

The mechanism described above is a highly efficient process that can handle the complexity of the many combinations and permutations resulting from the interaction of Drive Pools, Work Group Pools, JESplex resources, and asymmetric systems.

So the job selection process depends on two things:

- The DBS System Affinity of the job.
- The status of the Pools it needs.

The next two sections give a brief explanation about the processes of setting DBS System Affinity masks and the OPENing and CLOSing of Pools.

Managing the DBS System Affinity Assignment.

The initial DBS System Affinity mask for the job is constructed at job analysis time. It is based on the active Policy at the time of analysis. *If a job exceeds the resource requirements at the JESplex level for a Configuration,* the job is failed with an informative message and a JCL error. The requirements of the job can never be satisfied regardless of the Policy.

If a job exceeds the resources available to one or more systems with the active Policy, the DBS System Affinity mask is constructed to reflect the situation. The job is allowed to proceed to the execution queue.

When a new Policy is activated, the activation process reviews the Policy resources and the jobs in the queue and modifies their DBS System Affinity mask if necessary. In some cases, if the Policy is richer in resources affinities are added. In cases where the resources are more limited, affinities are removed.

OPENing AND CLOSing Pools

The decision to OPEN/CLOSE Pools is made by two separate processes:

- The DBS allocation algorithm. This is the most important and complex of the two processes.
- The step initiation process.

The DBS allocation algorithm is invoked at regular intervals. It can execute in any of the participating systems in a given JESplex. When it runs, it performs a JESplex wide evaluation. From this evaluation it determines what the new masks should be. These new masks are broadcast to all the systems using XCF. Each system thus has the most current JESplex/system mask in memory. The job selection process always uses the masks reflecting the current conditions.

At step initiation, if any request cannot be satisfied because there are not sufficient drives available, a signal is sent to have the “deficient” Pool CLOSED. It should be noted here that the step initiation process is not aware whether a particular Pool is OPENED or CLOSED. The step authorization process is based on the number of drives needed by the step and the number of drives that are available. The fact that a Pool is CLOSED does not mean that there are no drives available. The over-allocation algorithm might decide that there are enough jobs running with high watermarks that could result in over-commitment. In that case the Pool is CLOSED; however, the running steps will have their requests satisfied. It must be remembered that the point of the exercise is to run at the edge. That is, when a step request is made, there are just about enough avail- able drives to satisfy the request.

So to repeat:

- The fact that a Pool is CLOSED does not necessarily mean that there are no available drives. The DBS allocation algorithm might have determined that initiating another job that needs drives from that Pool will exceed the over-allocation factors. That is, the statistical possibility of causing an allocation recovery to occur is too high.

A Pool is always CLOSED when DBS has to make a step wait, however. It indicates that DBS *has gone beyond the edge* for the Pool so no more load can be accepted until the situation returns to normal.

DBS Unavailable Devices

Devices are unavailable (at a system or JlESpex level) when they can no longer be considered by DBS to be available for allocation to a Job. When a device is no longer available it will no longer contribute to the pool's available count. The maximum pool counts may need to be logically altered since they cannot be higher than the total number of actual devices in the pool.

This may affect the ability of jobs to run. There may be fewer or no systems where a job can run as the result of devices becoming unavailable. The jobs will not be able to run because there will be no system where the job has DBS affinity. The affected jobs will remain on the queue but will show a status of 'DBS affinity'. Examining the job or jobs with the DBS dialogue will show the details of 'why the job is not running'. By drilling down to the device level the displays will also indicate the reason the device(s) are unavailable.

- If jobs do not have DBS Affinity to any system, the DBS application will handle the jobs depending on the number of devices in the required pools. As long as there are enough devices to allow the possibility of the job running, the job will just remain on the queue but will not execute. The devices going into the count may be unavailable, they are still counted for the purpose of determining how to handle the job.
- If there are insufficient devices in the current configuration to satisfy the jobs requirements even when unavailable devices are added into the count then the job will be handled differently. If the job is already on the queue when this occurs then it will be put in MHS_HOLD for DBS. If a job hits this situation when it is analyzed then it will be put in MHS_HOLD and then requeued for analysis.

When device(s) in a pool are made unavailable with the command DBS SET DEV device#UNAVAILABLE some existing DBS jobs may be affected. A job that was selectable may no longer be selectable because it required a count that included the device that has been made UNAVAILABLE. These jobs will not be selected because the DBS affinity will now reflect the change in the available devices. They will NOT be held.

When enough devices are made available again the DBS affinity will be altered to make them selectable.

Chapter 4
Planning for DBS

This chapter includes topics for consideration when planning for DBS.

Introduction

DBS is designed to manage your tape drives with a minimum of intervention. Once the initial configuration definition is completed, DBS is normally transparent to both operators and users.

Before we can discuss the Configuration process we must introduce the Automation File.

The Automation File

DBS stores the Configuration and Policies in the *Automation File*. This file is like a Coupled Dataset for Automation Services, and is also used by SLM. It is shared among all systems in a JESplex. Installations with more than one JESplex will need *one Automation File for each JESplex*.

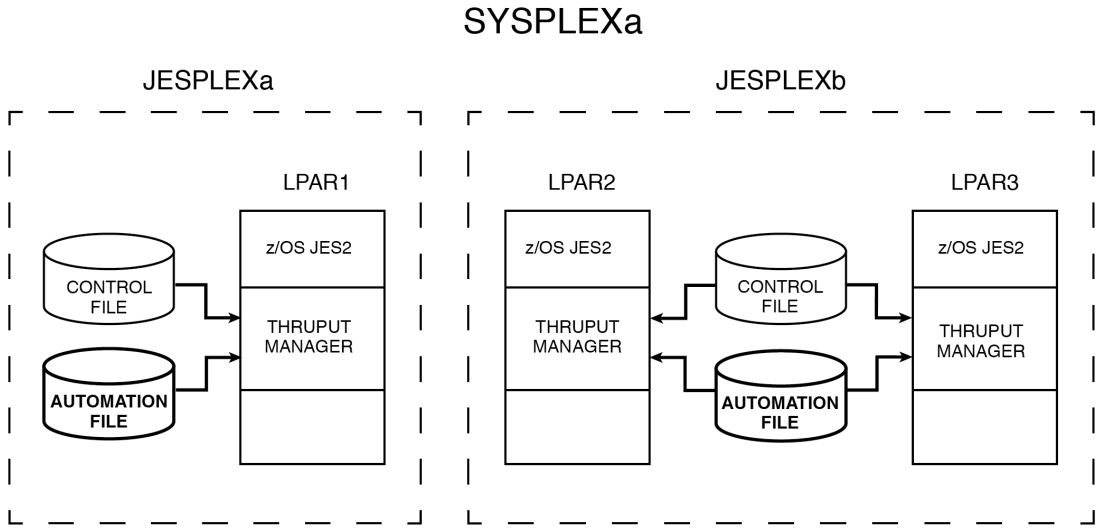

The Automation File is initialized and maintained using ISPF dialogs. When you define an Automation File for the first time, information describing the JESplex is collected. Since this description is used by other automation facilities, the Automation File for a particular JESplex needs to be created and initialized only once.

There are no performance considerations. You need only specify a volume that is shared across the JESplex. The ISPF dialog performs the actual allocation and initialization.

To activate a DBS Configuration from the Automation File:

1. The Control File must be converted to Version 6.
2. The DBS option must be enabled on the JES2 initialization statement TMPARM.
3. Use the DBS configuration dialog in ISPF to define your Configuration and any Policies you want to implement.

Configuration Definition

When referring to the creation of the initial or NEW configuration, we use the word definition as opposed to planning. The exercise is really a "fact-finding" process. The configuration in your installation is what it is, so you must collect the following information from your tape subsystem support group:

- The vendors that are present in your installation. DBS supports two hardware vendors and two software vendors:
  - IBM
  - StorageTek
  - CA-Vtape
  - EMC CopyCross
- The operational modes for each vendor:
  - Manual
  - Robotics
  - Virtual
- The names of all the libraries managing the devices. (The "names" for StorageTek robotics are predefined as ACS00 to ACSnn.)
- The way the actual transports have been SYSGENed. This is IBM's generic device type, for example 3490E.
- The actual device type and its description depending on its generic device type. See Supported Devices.
- The actual device numbers.
- Any asymmetries in the I/O configuration. For example, some transports might not be accessible to all the participating systems in the JESplex.

Without this information, there is no point in attempting to create the initial DBS configuration. The needed details should be readily available from the technical group that supports the z/OS I/O configuration.

The process of creating a Configuration is highly structured. Here we will discuss only the creation of the initial Configuration, which will be given the status of NEW. It maintains NEW status until it is activated. Once a NEW Configuration is activated, the status NEW is no longer used. When creating a future Configuration its status will be NEXT. The actual Configuration management cycle is described in the chapter Administering DBS.

It is important to review carefully the information gathered from the z/OS hardware support group. *You should include only the transports that DBS is to manage.* If you have transports that are not associated with traditional z/OS systems and as such are not available to allocation, they must be omitted from the definition.

The creation of a Configuration is the first step in the process. You can do this without having yet decided the Policies to be created. When the NEW Configuration is saved, a default Policy is automatically created under the name **BASE**. This Policy simply reflects the Configuration.

In your Configuration you can include "future" devices, so you can be prepared for future expansion (as long as it is known what the future is).

Note that DBS can handle transports that have different device numbers in different systems. It does *not* support environments with different devices having the same device numbers.

The Configuration is the framework that tells DBS what devices are to be managed.

Policy Planning

This is truly a planning process.

As already indicated, an initial default Policy named **BASE** is created automatically. This Policy is a one-to-one reflection of the Configuration. If your actual configuration is totally symmetric, then so is the default Policy. If the Configuration reflects hardware asymmetries, then so will the default Policy.

The **BASE** Policy can be edited to reflect whatever changes you might think are appropriate; however, since this Policy is your fall-back in case of problems, you should keep it as simple as possible.

The most obvious type of change you might want to make to the **BASE** Policy is the elimination of any transports that have been included to reflect "futures." Since they are not present, the actual counts reflecting availability should not include these future drives.

The construction of Policies requires a review of a number of considerations. The Configuration creation is a factual process. Policies, on the other hand, reflect deployment decisions. As such, a number of value judgments go into Policy creation.

We will defer a discussion of Work Group management at this time. This aspect of DBS requires management decisions with regard to resource allocation (who gets what). It is less a technical issue and more a business importance issue. A separate section below addresses the technical aspects of Work Group resource allocation.

In installations with a single JESplex and only one LPAR, the permutations and combinations on how to deploy tape transports are limited, so creating Policies should be simple. The creation of separate Policies might reflect situations where the installation wants to run restricted tape workloads during certain time of the day.

For example, there might be periods during the evening where the backup process for your distributed servers (assuming you back them up centrally) requires as many tape transports as possible. You might still want to run some critical batch work that requires tape. A Policy can address that need by simply restricting the number of drives available to batch during that period.

JESplexes with multiple LPARs present a number of alternatives to transport deployment under DBS management. A number of questions should therefore be considered, since it is possible to alter load distribution with simple Policy changes:

- If the hardware is symmetric:
  - Do you want to run an "equal opportunity" JESplex?
  - Do you want to bias the tape workload towards a particular LPAR?
- Do you want to direct certain type of work (for example, manual library requests) to only one system?
- Do you want to stop or minimize tape processing during a particular period of the day:
  - For the whole JESplex?
  - For a single LPAR?
- With Policies, you can make your hardware appear to be as asymmetric as you want.

  For example, you might have two LPARS and two Virtual Tape subsystems. During normal operations LPAR A has access to VT1 and LPAR B has access to VT2. You can create this asymmetry with a Policy (even though the hardware is accessible to both LPARs). DBS will select jobs only in the correct system. If one of the LPARs fails and the workload has to be accommodated in the other LPAR, a simple Policy change will do that.

- You might want to consider Policies for hardware failures, or for hardware service.

  For example, if you have multiple ACSs you can create appropriate Policies to pre-plan how to handle a particular ACS being out of service.

All Policies should have a name and description that reflect their purpose.

Work Group Planning

Because this area represents an allocation of resources to different workloads, it requires careful consideration. The mechanism provides significant flexibility to adjust resource allocation quickly. Activating a different Policy can drastically alter the tape resources available to a group.

Note that DBS Work Groups are optional.

What is a Work Group?

A Work Group is a collection of batch jobs that your installation considers to be related. The relationship is whatever you want it to be. It could be a similar level of importance even though the jobs come from different areas. It could be a type of work, such as production versus on demand. It could represent different divisions of the company.

The only requirement, from a DBS point of view, is that the jobs can be identified in JAL. The full power of JAL and DAL is available to identify the desired grouping. Once this is done, the job is assigned to the chosen Work Group.

Work Groups Versus Drive Pools

The Work Group concept is easy to understand, but if some simple rules are not followed it could be confusing to implement. Without it, DBS manages tape drives as z/OS does. There is no particular relationship between jobs and the available tape resources. All jobs are treated in a similar manner. Either there are drives (of a particular type) for everyone or for no one. There is no “best fit” algorithm or preference for jobs requiring fewer tape transports. The actual choice of which jobs is to be selected next is left to either the traditional JES2 class selection or to WLM Service Classes.

The role of DBS is to decide, once the job selection algorithm has made its choice, whether or not the job can be allowed to continue to initiation. At job selection time the decision is rather simple: either the pool(s) of drives needed by the job is opened or it is closed. If opened, the job is allowed to proceed; if closed, the job is bypassed.

The algorithm to determine when drive pools are opened or closed is complex and takes into consideration several factors. The approach is similar for Drive Pools and Work Group Pools. There is, however, an essential difference that is important to understand before values for Work Groups are assigned.

Drive Pools are defined at the lowest possible level of device groupings. For example, if you have an IBM automated library, let’s say TAPELIB1, with 8 3490s and 16 3590s, then DBS constructs two Drive Pools. One pool represents the 8 3490s, the other pool the 16 3590s.

If your installation were to add another similar library, TAPELIB2, then DBS constructs 4 Drive Pools. For Drive Pools, drives that are part of TAPELIB1 are not considered to be part of the same pool as drives from TAPELIB2. They are not interchangeable, even though they are the same type of drives. This has to do with volume accessibility.

For the purpose of defining values for Work Groups, the aggregation is different. When designing DBS it was decided that managing them at the l*owest level* was impractical: too much detail and too much segmentation. It would have made the task unnecessarily complicated. So the decision was to define Work Groups at the *Vendor/Mode/level*.

In the example described above, the definitions for Work Groups will be at the IBM/ AUTOMATED mode. Only two Work Group pools are created regardless of the number of similar (automated) libraries that are defined. So, from a Work Group definition point of view, you are given the opportunity to apportion drives from two Work Pools:

```
IBM/AUTOMATED/3490/3490E     16 Drives
IBM/AUTOMATED/3590/3590E     32 Drives
```

To repeat:

- Drive Pools are defined at the lowest possible level.
- Work Group Pools are defined at a higher level: VENDOR/MODE level.

The ISPF dialog automatically aggregates the counts, so you know what the total numbers are at any level.

The mechanism allows you to define a value that represents the maximum number of drives (of a given vendor/mode/type) that a particular Work Group can have allocated at any given time.

Apportioning Drives Using Work Groups

The second and sometimes confusing consideration is the process of apportioning the drives. This must be done as a separate process and in agreement with whatever group or groups need to be consulted in your installation. This differs from place to place because of organizational arrangements. You can apportion drives in different manners. What follows is a discussion of the facilities provided by the Work Group option.

You can dedicate drives to a particular Work Group. As the name indicates the drives (a count, not specific device numbers) are for the exclusive use of that Work Group. The dedicated number of drives—let's say 5—are not available to any other Work Group even when they are idling. Again, what is dedicated is a number of transports, not devices such as 0F01, 0F02, etc. If we use an airline booking analogy, DBS reserves the appropriate number of "tickets" with no specific seating arrangement. Standard allocation when a "ticket" holder arrives takes care of the actual "seating arrangement."

Of course, the total number of dedicated drives cannot exceed the number of available drives. Normally, the number of dedicated drives should be smaller than the number of available drives, so the difference between available drives minus dedicated drives represents the number of non-dedicated drives.

By default, all Work Groups contend for access to the non-dedicated drives: first come, first served. At any given time there might not be any available; at another time, a particular Work Group could be hoarding most if not all of the non-dedicated drives. You can handle the problem described above by controlling the maximum number of drives available to a particular Work Group. This is known in DBS as CAPPING a Work Group.

When a Work Group has been CAPPED, it does not matter whether or not additional drives are available and are idle. The CAP is a hard restriction.

The value associated with a CAP represents the sum total of the dedicated drives plus the "get-them-if-you-can" non-dedicated drives. So, assume that the values assigned to a Work Group are as follows:

```
Dedicated    6
CAP          12
```

This means:

- The Work Group will always have at its disposal 6 drives.
- In addition, it will be able to contend for another additional 6 drives.

The Work Group dialog performs all the necessary calculations. As a result of 6 drives being dedicated, the number of non-dedicated drives available is reduced by 6. Of course, it also ensures that there are, at least, 6 non-dedicated drives.

The dialog is very fluid because as soon as you dedicate drives the number of non-dedicated drives change.

Implementing Work Groups

What the dialog cannot do for your installation is decide how to create effective Work Groups and the apportion drives across them. To implement Work Groups, you have to do the following things:

- Determine whether you want to implement Work Groups at the JESplex level or at the individual JESplex member level.
- Decide how jobs are to be grouped. It could be as simple as production versus non-production.

- Based on the requirements resulting from the groupings, apportion drives to balance the different requirements.
- Consider constructing more than one Policy if needs change during the day. Your production workload might need most of the drives during some critical periods. This is easily done with different Policies.

Before considering the segmentation of your work load into different Work Groups, you should clearly understand that Work Groups are defined at two interdependent levels:

- The higher level allows for three partitions.
- Each high level partition can be further subdivided into two.
- The total number of Work Groups is therefore six (3 x 2).

The DBS Work Group structure is represented in the diagram below.

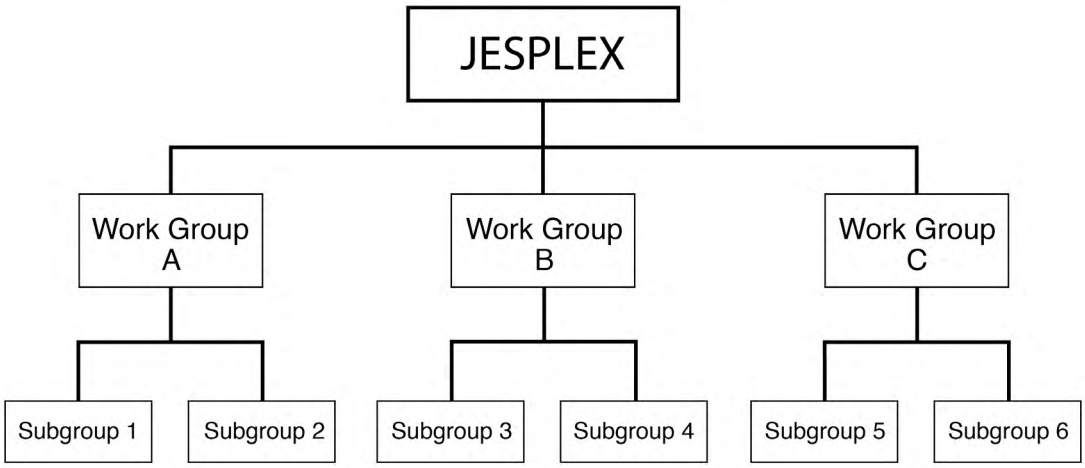

The actual apportioning of the available tape transports is done at the highest level. That is, they can be partitioned into three mutually exclusive segments. The first level can then apportion whatever drives have been apportioned to it from the general pool to its two Work Groups.

The first level partition represents a *global* decision, where all tape users are affected.

The second level partition is a *local* decision, because no additional drives can be apportioned that have not already been given to the first level. If no dedicated drives were given to the first level, no dedicated drives can be given to the second level. If the first level was CAPPED, the value of the CAP is automatically reflected in the second level.

In summary:

- Dedicating drives eliminates the possibility of a Work Group ending up with no drives.
- CAPPING eliminates the possibility of a Work Group hoarding all the non-dedicated drives.
- The overall allocation of drives is done at the first level of Work Groups.
- The second level of Work Groups is restricted by the values assigned to its corresponding first level.
- Drives that are dedicated to the first level do not necessarily have to be dedicated to the corresponding second level.
- All Work Groups are given unique names (installation chosen).
- When assigning a Work Group to jobs in JAL, the name of the second level is used.
- DBS displays show the name of the first and second level.
- Work Groups are associated with a Policy, not a Configuration.

The Definition Process: What Should I Have?

Before starting the DBS Configuration and Policy definition process you should have information that is similar to the following example:

- The configuration has to be named. Let's call it CONFIG01.
- The JES2 node name is needed: CAMPUS1.
- For verification purposes, the SPOOL dataset name is requested. Let's say SYS1. HASPACE.
- The SPOOL Volume prefix: CAMP1 in this case.
- The Participant JES2 member names. In this installation there are two LPARS: SYS1 and SYS2.
- The vendor or vendors, in this case IBM.
- The Modes, in this case Manual and Virtual.
- The high level devices: for Manual, 3590-1; for Virtual, 3490V.
- The actual devices under the high level: the 3590-1 are 3590E, the 3490V are 3490E.
- The VTS Library name, VTS01 for this example.
- Whether you plan to use Work Groups: NO initially.

Now the device counts and device numbers are needed:

- 6 3590E devices:
  - Devices numbers 0360-0365.
- 64 Virtual 3490E devices:
  - Device numbers 0380-03BF.

The above information is all that is needed for DBS to do its basic job. We will start with the **BASE** Policy.

To sum up:

```
CONFIG NAME           CONFIG01
JESPLEX NAME:         CAMPUS1
SPOOL DSNAME:         SYS1.HASPACE
PREFIX:               CAMP1
JES2 MEMBER NAMES:    SYS1, SYS2
VENDOR:               IBM
  MODE: MANUAL
    DEVICE TYPE:      3590-1
    SUBTYPE:          3590E
    DEVICE COUNT:        6
    DEVICE NUMBERS:   0360-0365
  MODE:               VIRTUAL
    DEVICE TYPE:      3490V
    SUBTYPE:          3490E
    DEVICE COUNT:       64
    DEVICE NUMBERS:   0380-03BF
```

With the above information you can create the NEW Configuration. The dialog will automatically create the **BASE** Policy.

Let's introduce Work Groups. You could do that by editing the **BASE** Policy, but we do not recommend that approach. The **BASE** Policy should be kept as simple as possible as a fall-back in case of problems. A new Policy should be created.

For the purpose of this example, let's assume that the Work Group definitions are symmetric. (It is possible to have different values for each system.) Further, let's assume that Work Groups are not needed to manage the Manual devices.

So, for the virtual devices we want to create Work Groups. In order to do so we have to do some fiction writing.

- This installation wants to divide its batch workload into three FIRST LEVEL Work Groups:
  - PROD for production jobs
  - NON_PROD for non-production jobs
  - DEV for the development group

- The PROD Work Group is to have two sub-groups:
  - SPECIAL
  - NORMAL
- The NON_PROD Work Group is to have two sub-groups:
  - TYPE1
  - TYPE2
- Again, the DEV group has two sub-groups:
  - DEV1
  - DEV2

After the usual endless meetings, the drives (64 of them) are allocated as follows:

- PROD has access to as many as 48 drives. Of the 48 drives, 16 are dedicated. 16 drives are not available to PROD.
- NON_PROD gets to access up to 32 drives. They are given 8 dedicated drives.
- DEV has access up to 24 drives. They are given 6 dedicated drives.

Now we can examine how to express these values to DBS. For LIBRARY 1, named VTS01, and for 3490E devices, the following Work Groups are defined:

```
WORK GROUP 1
  NAME:           PROD
  DEDICATED:      16
  CAPPED:         48
WORK GROUP 2
  NAME:           NON_PROD
  DEDICATED:      8
  CAPPED:         32
WORK GROUP 3
  NAME:           DEV
  DEDICATED:      6
  CAPPED:         24
```

You can now divide each Work Group into two subgroups. Here the decisions are no longer “global.” The Groups are apportioning what is already theirs. As a result the meetings should be smaller:

- Let’s assume that the PROD group wants to dedicate 4 drives for some special jobs. The rest is all contention among “normal” jobs.
- In the case of NON_PROD they want to dedicate 6 drives to one subgroup. All the other drives are available to both sub-groups.
- Finally, the DEV group wants to do a fair split of drives for the two subgroups but ensure that they cannot be monopolized by either group.

What follows is the implementation of the above rules. They probably represent several days of discussion, but with the magic of DBS they can be implemented in minutes.

The complete Work Group Definitions, if they were to be entered in a table as opposed to a dialog, would look like the following:

```
WORK GROUP 1
  NAME:  PROD        DEDICATED:   16  CAPPED:  48
   SUBGROUP 1
     NAME:  SPECIAL  DEDICATED:    4  CAPPED:   4
   SUBGROUP 2
     NAME:  NORMAL   DEDICATED:   12  CAPPED:  44
WORK GROUP 2
  NAME:  NON_PROD    DEDICATED:    8  CAPPED:  32
   SUBGROUP 1
     NAME:  TYPE1    DEDICATED:    6  CAPPED:  32
   SUBGROUP 2
     NAME:  TYPE2    DEDICATED:    0  CAPPED:  26
WORK GROUP 3
  NAME:  DEV         DEDICATED:    6  CAPPED:  24
   SUBGROUP 1
     NAME:    DEV1   DEDICATED:    3  CAPPED:  21
   SUBGROUP 2
     NAME:   DEV2    DEDICATED:    3  CAPPED:  21
```

Is it obvious what the numbers mean? Well, they may require a bit of an explanation since each Work Group affects the others:

- The maximum (CAPPED) that Work Group PROD can ever allocate is 48 drives.
  - If that were to occur then there are only 16 drives left for NON_PROD and DEV.
  - Of the 16 drives, 14 are already preallocated because NON_PROD has a DEDICATED count of 8 and DEV has a DEDICATED count of 6.
  - So, in this situation there are only 2 drives for NON_PROD and DEV to contend for.
- The DEDICATED count indicates that regardless of the availability of work for that Work Group, the drives are reserved and not accessible to other Work Groups.

  For example, even if no jobs for Work Groups PROD and NON_PROD are in the JESplex, the maximum number of drives available to DEV is 40 even if no restrictions were placed on DEV. As defined, the Work Group is restricted to 24.

- The CAPPED number of drives includes the "dedicated drives" (DEDICATED), which could be zero, and the shared drives which could be zero also.

- The sum total of the DEDICATED values for SUBGROUPS cannot exceed the DEDICATED value for the Work Group to which they belong. This means that you cannot dedicate drives to your SUBGROUPS that have not been dedicated to the Work Group. The reverse is not necessarily true. You do not have to dedicate drives to the second level just because the first level has dedicated drives. You can allow the corresponding two Work Groups in the second level to have an equal chance to get any of the dedicated drives. The Work Group PROD/NORMAL could be defined as DEDICATED 16, CAPPED 48 or DEDICATED 0, CAPPED 48 with equal results. This is a consequence of the way PROD/SPEC is defined.

Note that the calculations are done automatically by the Work Groups definition dialog. For example, if there are 64 drives in the Drive Pool and you assign a DEDICATED count of 24 to a Work Group, the number of drives available for apportioning to other Work Groups is reduced to 40 automatically.

Now we can show why you can define multiple Policies to handle particular situations. Let's say that the situation described above represents *daytime* drive allocations for Work Groups. This installation re-apportions the tape drives differently after midnight to reflect the needs of the production cycle. At that time they apportion most of the drives to the production group. The development group does not have any drives during that period. This Policy might be called NIGHT:

```
POLICY NAME:        NIGHT
WORK GROUP 1
   NAME: PROD       DEDICATED:  40  CAPPED:  56
    SUBGROUP 1
      NAME: SPEC    DEDICATED:   4  CAPPED:   4
    SUBGROUP 2
      NAME: NORMAL  DEDICATED:  36  CAPPED:  52
WORK GROUP 2
   NAME: NON_PROD   DEDICATED:   0  CAPPED:  18
    SUBGROUP 1
      NAME: TYPE1   DEDICATED:   0  CAPPED:  18
    SUBGROUP 2
      NAME: TYPE2   DEDICATED:   0  CAPPED:  18
WORK GROUP 3
   NAME: DEV        DEDICATED:   0  CAPPED:   0
    SUBGROUP 1
      NAME: DEV1    DEDICATED:   0  CAPPED:   0
    SUBGROUP 2
      NAME: DEV2    DEDICATED:   0  CAPPED:   0
```

By simply activating the "night" Policy, the Production group has control of most of the drives.

Another Policy can be created in case the Production Group runs into serious difficulties. The EMERGNCY Policy might look like the following:

```
POLICY NAME:         EMERGNCY
WORK GROUP 1
   NAME: PROD        DEDICATED:  4  CAPPED:  64
    SUBGROUP 1
      NAME: SPEC     DEDICATED:  4  CAPPED:   4
    SUBGROUP 2
      NAME: NORMAL   DEDICATED:  0  CAPPED:  60
WORK GROUP 2
   NAME: NON_PROD    DEDICATED:  0  CAPPED:   0
    SUBGROUP 1
      NAME: TYPE1    DEDICATED:  0  CAPPED:   0
    SUBGROUP 2
      NAME: TYPE2    DEDICATED:  0  CAPPED:   0
WORK GROUP 3
   NAME: DEV         DEDICATED:  0  CAPPED:   0
    SUBGROUP 1
      NAME: DEV1     DEDICATED:  0  CAPPED:   0
    SUBGROUP 2
      NAME: DEV2     DEDICATED:  0  CAPPED:   0
```

If this Policy is activated, then NON_PROD and DEV groups should go fishing.

DBS and Job Limiting (JLS)

Prior to the availability of DBS, the Job Limiting Services facility of ThruPut Manager has been the mechanism for controlling tape usage. DBS has largely taken over this function, but for some installations, JLS still has a role to play. This section explains how JLS can be used to extend the control over tape device apportioning even beyond the level provided by DBS Work Groups.

The Old Role of JLS

JLS is a general purpose facility. As such, it is a more abstract mechanism than DBS. DBS "knows" that its mission is to help with the process of tape drive management; JLS does not. DBS deals with actual device counts, while JLS deals with abstract counts. As a result, the installation has to make the connection between the abstract counts and the actual number of tape drives available. ThruPut Manager assumed the responsibility of establishing high watermarks for the job and, if requested, reducing them at step termination if the step in question was the one reflecting the current high watermark.

The job selection mechanism of ThruPut Manager determined, among other things, if there were enough units (in the abstract count) to accommodate the actual count associated with the job. If not, the job was not selected.

Since the number associated with the job represented the high watermark, underutilization of drives occurred unless the installation overbooked the tape drives by defining an abstract number higher than the actual number.

The New Role of JLS

None of these considerations are pertinent when DBS assumes the tape drive management. DBS has overbooking algorithms built in. They apply to general Drive Pools as well as values associated with Work Group Pools. You might, however, have situations where you want to restrict access to drives beyond the segmentation provided by DBS Work Groups. In essence, you want a finer degree of granularity. This has to be done with JLS.

To facilitate this process, DBS introduces a set of JAL Descriptors that reflect the values used for the management of DBS Work Groups. These new Descriptors represent the high watermark of a particular DBS Pool. Following normal JLS rules, the Descriptor can be associated as "weight" for a given Limiting Agent, which in turn is associated with the job.

The high watermark value for these Descriptors is adjusted (downwards) as required. This occurs automatically for any Limiting Agent that uses the new DBS Descriptors.

An Example of Using JLS to Extend DBS

Let's discuss an example where Work Groups are extended by the use of JLS.

One of the Work Groups, which we will call DEV2, represents an IS development group. The individual developers do not have jobs that require more than 2 tape drives. In most cases, one drive is all that they require for an individual job; however, it is possible for a developer to submit multiple jobs at once. This prevents other members of the group from having access to drives until the jobs from that particular developer terminate.

When the general apportioning of drives to Work Groups took place, DEV2 was given a maximum of six drives. Within the facilities provided by DBS, there is no way of creating further restrictions to control individuals, even though the members of the group are prepared to accept a restriction of two concurrent drives per individual. Without the implementation that automatically controls the agreement, the only way is to manage the jobs by submitting them in HOLD.

Here is where JLS can play a role by complementing the Work Group segmentation of DBS. Let's say you can identify individual developers by User ID. A Limiting Agent for each developer can be created dynamically in JAL with a maximum value of 2, and a "weight" represented by the appropriate DBS Descriptor. That is all that is needed to have a totally automated solution:

- DBS manages the Global Pool.
- DBS enforces the Work Group values (6 in this case).
- JLS enforces the individual maximum (2 in this case).

The structure for using JLS to extend DBS Work Groups will look like this:

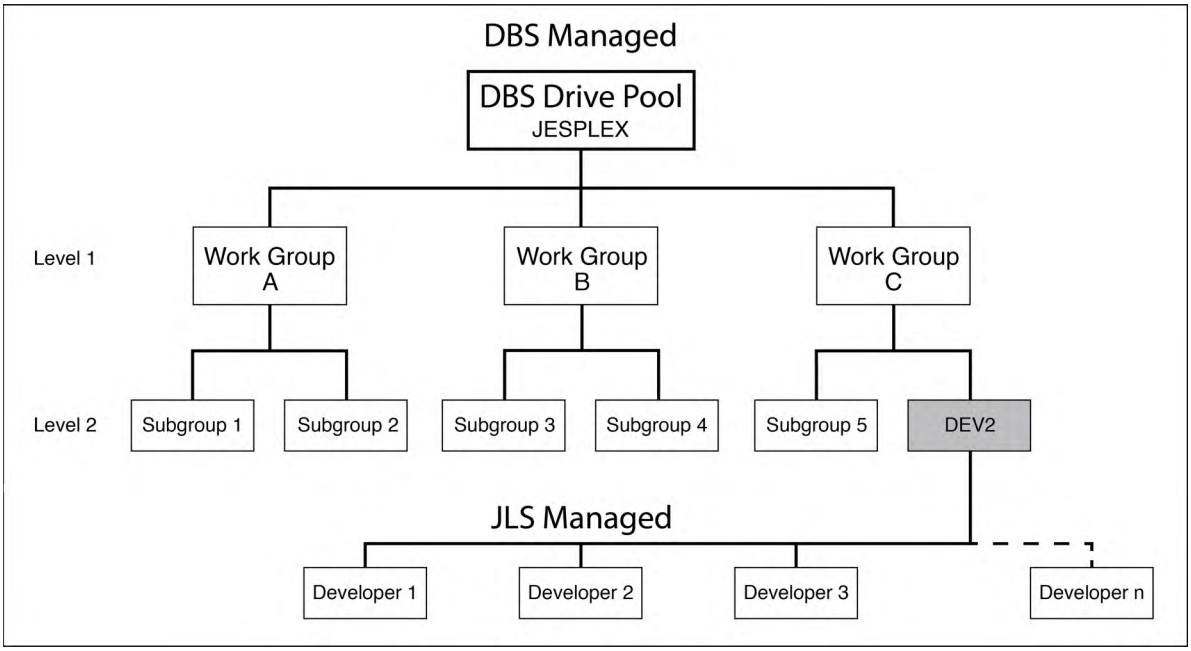

When combining DBS and JLS to control tape drive usage remember the following rules:

- Do the overall tape usage control with DBS.
- Do the next level of segmentation with DBS Work Groups.
- If you need finer granularity use JLS with the new DBS Descriptors.

Always subordinate JLS to DBS to get the best results.

Relating DBS Descriptors to $UNIT Descriptors

The new DBS Descriptors and their relationship to the $UNIT Descriptors are shown in the *Table of Supported Drive Pools* accompanying the description of the $DBS_ drivepool Descriptors in the *JAL Reference Guide*. In some cases, there is not a one-to-one relationship. If there were, there would not be a need for new Descriptors. The most obvious difference is the DBS separation of Descriptors by vendor. Installations with only one vendor will see fewer differences than installations with multiple vendors.

Chapter 5
Implementing DBS

This chapter describes the steps and procedures required to Implement DBS.

System Requirements

In order to implement and activate DBS:

- You must be running Version 7 of ThruPut Manager.
- The Control File must be in Version 7 format. If you are running ThruPut Manager Version 5, this requires a Control File conversion.

The scope of DBS is the same as that of your JES2 node.

DBS Facilities

The facilities provided to support DBS are summarized here.

Installation and Implementation

To install and implement DBS, you must:

1. Convert your Control File to Version 7. Although you can install the DBS modules without converting the Control File, the application cannot run, and a warning message will be issued.
2. Install the DBS modules as per the instructions contained in the accompanying Installer's Guide.
3. Use the TMISPF command to initiate the TM Automation Services Dialog, then select TM Automation File Services from the Main Lobby to begin defining a DBS Configuration. See DBS Configuration Definition Considerations.
4. When DBS is installed according to the Installer's Guide, the option name DBS is added to the selection of options that can be controlled through the OPTIONS keyword of the JES2 initialization statement TMPARM and the TM OPTIONS operator command. The default for this option is disabled. Change the JES2 initialization statement TMPARM to enable DBS.
5. Optionally, add the DBS keyword to the TM SMF initialization statement.
6. Stop TMSS on all systems.
7. Hotstart JES2. To enable DBS temporarily without a hotstart, you can use the operator command TM OPTIONS ENABLE DBS.
8. Restart ThruPut Manager on all systems.
9. Invoke the TM Automation Services dialog and activate the NEW Configuration.

After these steps have been completed, DBS is installed and ready to manage your tape devices.

DBS Configuration Definition Considerations

The DBS Configuration is kept in the Automation File, which is managed with ISPF dialogs.

Invoking the DBS Configuration Dialog

To define a DBS Configuration, first issue the TSO command under ISPF to invoke ThruPut Manager ISPF services:

```
TMISPF
```

This takes you to the Main Lobby for ThruPut Manager dialogs. From the list, select the ThruPut Manager Automation Services dialog. This brings up the Primary Options panel, from which you select TM Automation Services. You are now presented with a list of Automation File management options.

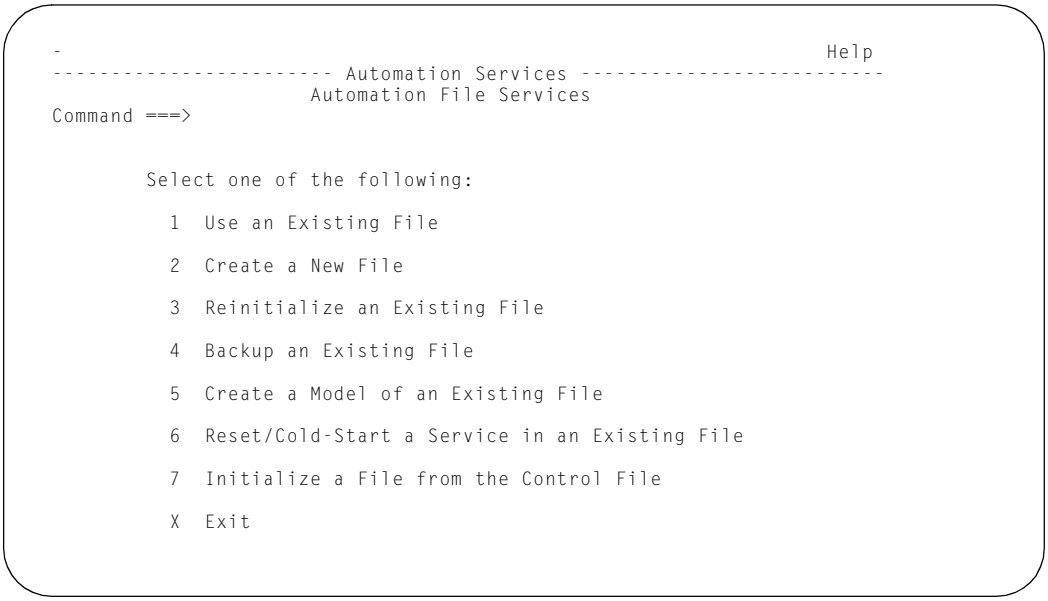

The Automation File is used by other ThruPut Manager services besides DBS, but it needs to be allocated and initialized only once. For our example, we assume that the file does not exist. Select the option to allocate and initialize a new file. This opens a window asking for details of the new Automation File.

Note that you need one Automation File for each JESplex.

JESplex Configuration

After you have provided a name for the Automation File, you are prompted for a descrip- tion of the JESplex. If you are running on the JESplex that will use the new file, you can have DBS automatically gather the information for you, otherwise you must provide it manually.

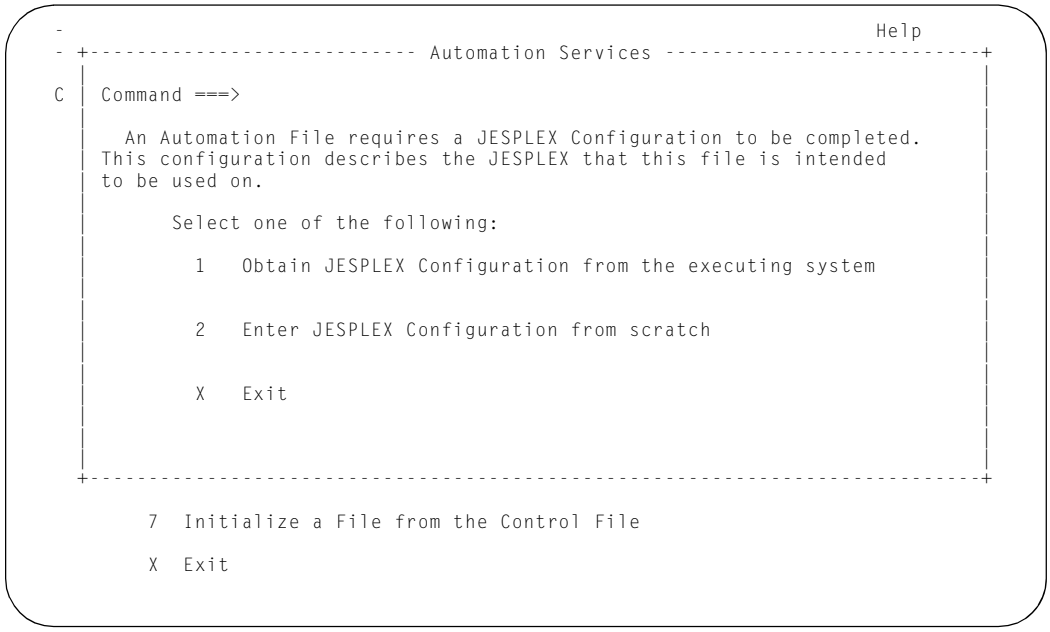

This JESplex description is common to all services that use the Automation File, including DBS. It includes:

- A description that will help identify the Automation File.
- The JES2 node name.
- The JES2 SPOOL dataset name.
- The JES2 SPOOL volume prefix, required to ensure that the JESplex is uniquely identified.

All the information provided here can be edited later if necessary.

Defining the First Configuration

Once you have saved the JESplex information, you should select the DBS service and define a DBS Configuration.

A description of the steps to follow to define a Configuration is provided in the chapter Managing the Configuration

The **BASE** Policy

When you save your first Configuration definition, DBS automatically creates a default Policy for you. A pop-up window prompts you to enter a description. The Policy has these characteristics:

- The Policy name, which cannot be altered, is “**BASE**”.
- All device counts are equal to the number of Device Numbers that youspecified in your Configuration.
- The Work Group feature is turned OFF.

You are placed in Policy Edit mode, which allows you to manage the Device Counts for the Drive Pools defined in the Configuration, and to turn on and manage the Work Group feature.

Although you can edit the **BASE** Policy at this point, we recommend that you do not make this Policy a specialized one for these reasons:

- This is the Policy that an operator can always restore with a simple command.
- It is a good idea to maintain a stable **BASE** Policy that you can use as the starting point for creating all other Policies.

For a description of creating specialized Policies, see the section Managing Policies.

DBS Work Groups

DBS Work Groups are helpful in distributing tape devices among various types of work. For example, you might want to make sure that test jobs don't exhaust a Drive Pool at the expense of production work. DBS allows you to do this by using JAL (or JECL) to assign jobs to different Work Groups. Each Work Group has limits on the minimum and maximum count of devices that it can allocate from the Drive Pool. Work Groups therefore determine "who gets what."

The initial Configuration definition does not turn on the Work Groups feature. We recommend leaving Work Groups turned off until you are satisfied that you have defined your Configuration and Policies correctly, and that DBS is running smoothly. Work Groups are not mandatory, and you might find that your installation does not need them at all.

When defining the Work Groups, keep the following in mind:

- Choose descriptive names. This will help later when viewing displays or analyzing results.
- There are three possible Work Groups, each with two Subgroups. You do not have to use all of them, but you cannot add new ones.
- The three Work Groups are for organization purposes. You actually assign one of the six Subgroups to jobs.
- The Subgroups have an implied priority. *All other things being equal*, the Subgroup priority ranking is from Subgroup 1 to Subgroup 6, or left to right as seen in a display of Work Groups. This order cannot be changed.

SMF Monitoring of DBS Activity

A ThruPut Manager facility is provided for an installation to gather statistical information about the activity of DBS. An SMF record containing this information is generated at intervals of approximately every ten minutes.

Activating Data Collection for DBS

To collect data from the DBS application, you must indicate this on the TM SMF initialization statement:

```
TM SMF TYPE(255) ANALYZER(9,SMFDATA) DBS
```

Identifying DBS Records

DBS records are identified by their subtype of 240. The subtype is found in the field named SMFTMSTP in the DTM DBSMF mapping macro. Refer to this macro for details of the collected data.

Chapter 6
Administering DBS

This chapter describes how to manage the DBS Configuration and Policies, and how to use DAL and JAL with DBS.

Introduction to DBS Administration

Once you have defined a DBS Configuration, there is no reason to change it unless your installation adds or removes tape devices. If you have included future devices in your definition (as we recommend) then the need for change is reduced even further.

Policies, on the other hand, provide flexibility. Using Policies, you can adjust DBS to adapt to the prevailing conditions. For example, Policies can reflect the difference in workload profiles during the day and overnight, or on weekends. Policies can be used to support contingency plans, such as device failures or planned maintenance outages. In short, it is likely that you will want to define several Policies. You probably do not have to change Policy definitions often, but you will want to switch Policies to meet the changing needs of the installation.

Managing the Configuration

A DBS Configuration can have only one of the following states:

- NEW: This represents the first Configuration before it is Activated.
- ACTIVE: This represents the Configuration that is in use by DBS for a particular JESplex. *The ACTIVE Configuration cannot be edited.*
- BACKOUT: This represents more than just the previous ACTIVE Configuration. The BACKOUT Configuration is a bridge between the previous ACTIVE Configuration and the running ACTIVE Configuration. This allows a fallback to a BACKOUT that takes into account jobs in the queue (or executing) that were processed with the NEXT Configuration but have to be adjusted to the BACKOUT.
- NEXT: This represents the future Configuration, as created by the CN command from the Configuration Management dialog.

Configuration Definition

To define a DBS Configuration, select the DBS service from the TM Automation Services menu to invoke Configuration Management. If this is your first Configuration, your only choice is to Create a NEW Configuration; otherwise, you can create the NEXT Configuration. In both cases, the command is CN.

Vendor Selection

The next step is to identify the vendors for the devices in your installation. DBS provides a structure that includes all the vendors, device types, and modes of operation that are supported. Select the appropriate vendors with the **S** line command.

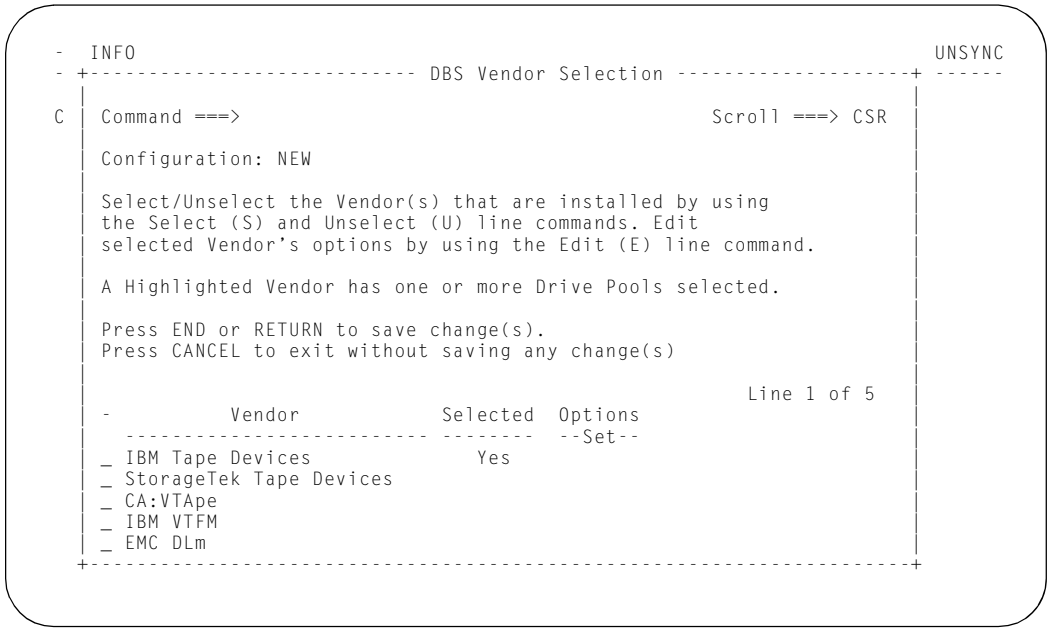

The vendor IBM is provided as a default. If you do not have IBM devices, you can use the **U** line command to deselect this vendor.

Selecting a Mode and Naming a Library

After you have selected the vendors, the Drive Pool/Device Number Management screen appears, displaying a tree structure for the tape drives..

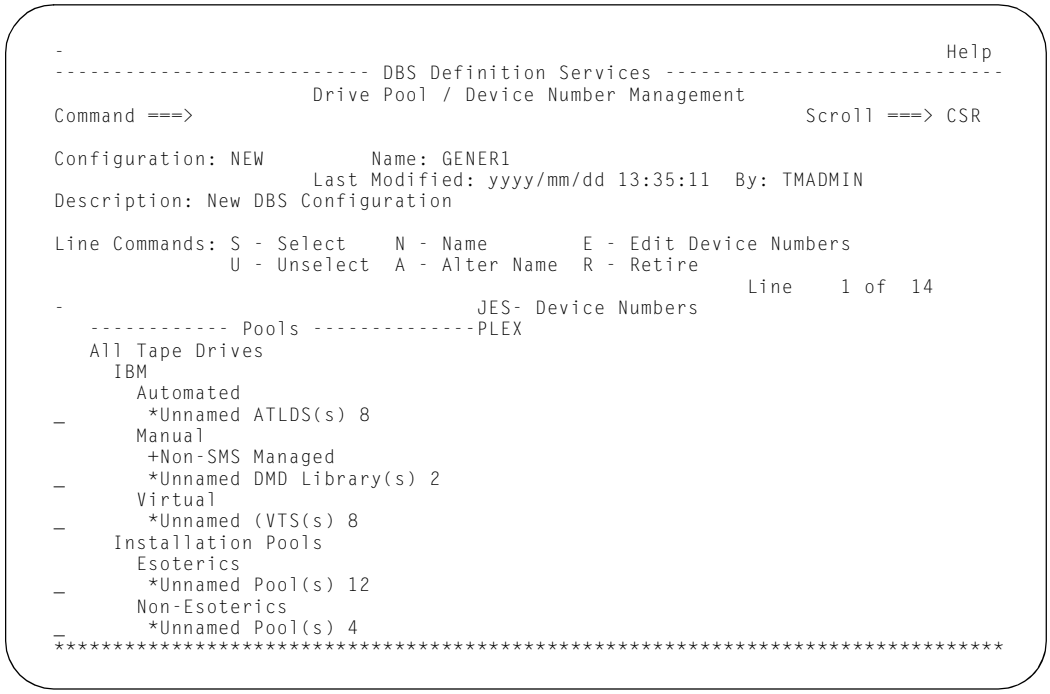

DBS needs to know the modes that you are using for your devices and, for modes that use them, the names of your libraries. The available modes are:

- IBM
  - Automated
  - Manual
  - Virtual
- StorageTek
  - Automated
  - Non-library
  - Virtual
- CA:Vtape
  - Virtual
- IBM VTFM
  - Virtual
- EMC DLm
  - Virtual

Only Pool names that begin with an asterisk (*) need to be named. The N line command invokes a pop-up window that allows you to do this, but first you must expand the tree structure so that you can choose the type of library that you want to name. A plus sign (+) indicates that a branch is collapsed, and a minus sign (-) indicates that the branch is expanded. You expand or collapse branches by placing the cursor within the Pool name and hitting ENTER. The N line command invokes a pop-up window that allows you to provide a library name.

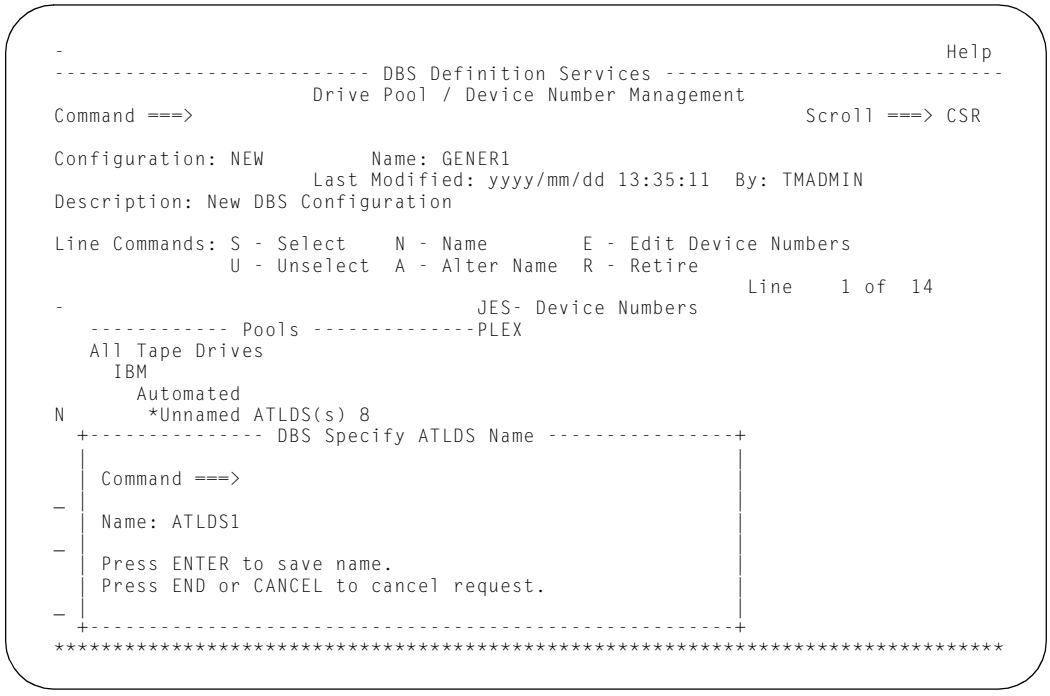

The form for the library name varies depending on the vendor:

- For IBM, enter the 1-8 character name matching the SMS-managed library you are defining.
- For StorageTek, choose the appropriate ACS name from the list.
- For CA-Vtape and EMC DLm, you do not need to specify a name.
- For Installation Pools (described further in Using Installation Pools) enter a name of your choosing. It must be 1-8 alphanumeric or underscore characters, the first of which must be alphabetic.

Note that by implication, the named library also specifies the device mode.

Selecting Device Types

Once a library has been named, the possible devices (generic type) are shown. This could require you to expand the branch for the device type you want to select. Expand if necessary, then use the S line command to select the generic device type under which the actual drives have been grouped.

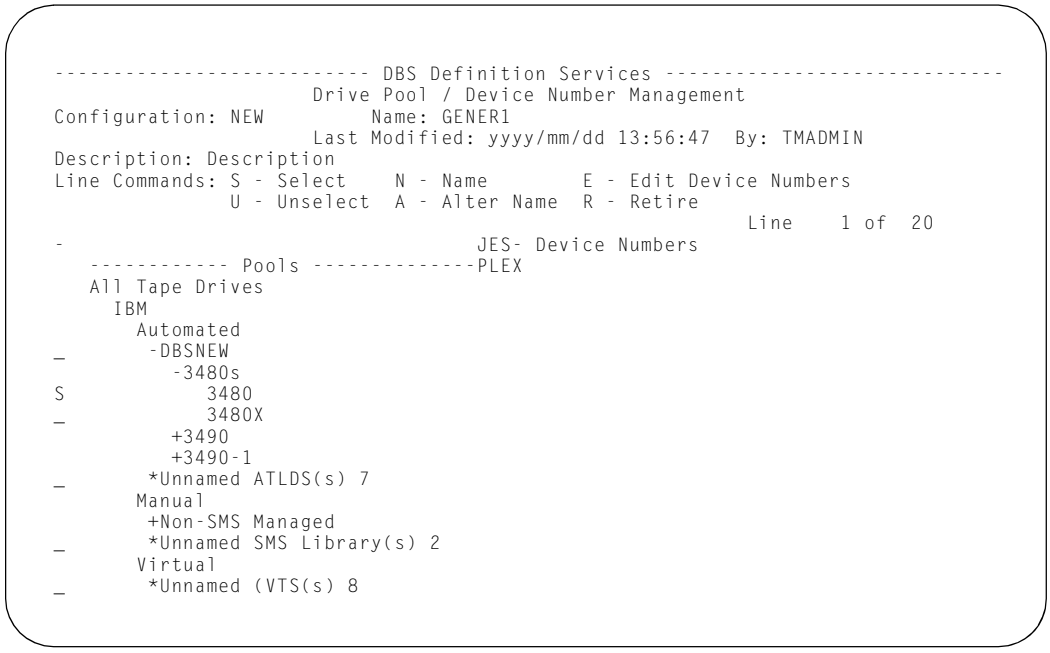

Specifying Device Numbers

The next step requires you to identify which Device Numbers belong to the selected Drive Pool.

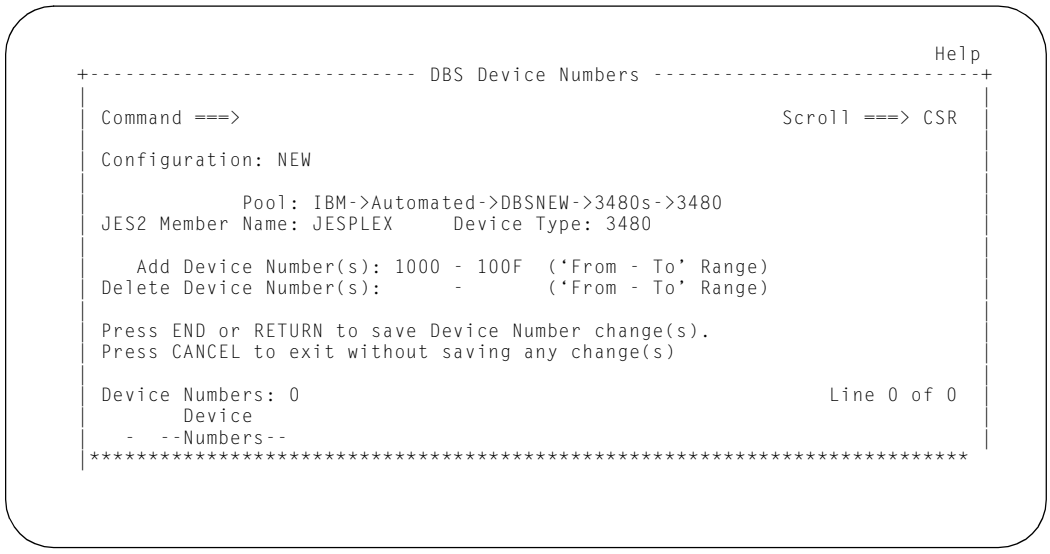

When you select the device type, the DBS Device Numbers panel opens, allowing you to specify the range(s) of *all Device Numbers available in your JESplex*.

Of course, it is possible that the members in your JESplex are asymmetric, i.e. some devices are not available on all systems. To reflect this, you can use the **E** line command to edit the Drive Pool after you have saved the ranges available for the entire JESplex. This sequence ensures that you do not inadvertently define Device Numbers for a JESplex member that are not already defined for the whole JESplex.

Note that the way you handle asymmetric devices depends on the type of asymmetry:

- If the devices are *physically* unavailable to some JESplex members, use the **E** command to reflect this in the Configuration.
- If the devices are *logically* unavailable to some JESplex members, use a Policy to reflect this. See Setting Device Counts."

Normally, a device number can be assigned to one Drive Pool only. An exception occurs when you elect to use Installation Drive Pools. For more information, refer to Using Installation Pools.

Repeat the sequence of naming libraries (where required) and selecting generic device types until you describe the entire configuration for your JESplex. Upon completion, save the Configuration to finish the definition process.

Activating a Configuration

To implement a Configuration, you must use the AN command of the Configuration Management dialog to activate it:

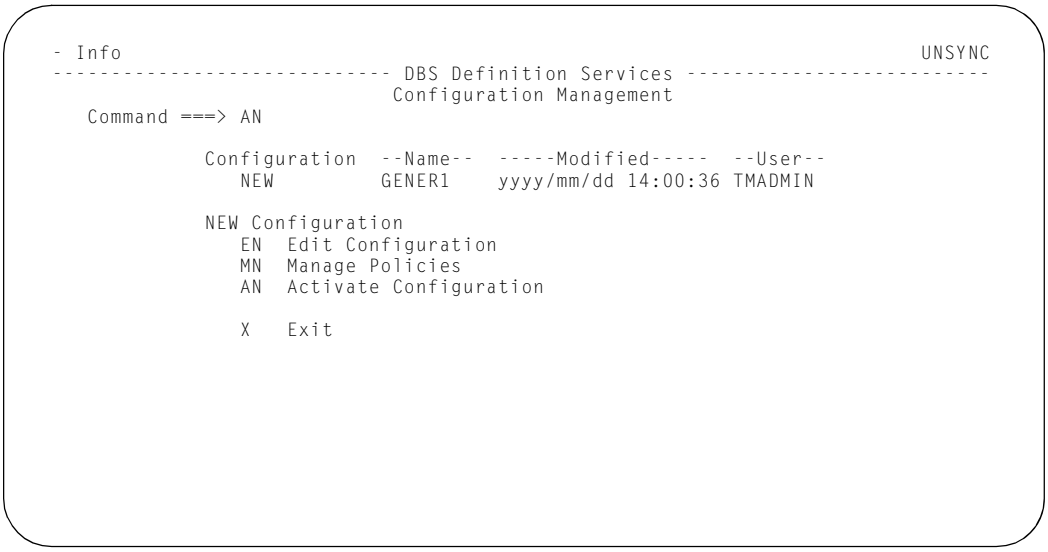

The procedure is identical for both NEW and NEXT Configurations. The difference is that a NEW Configuration has never been activated, while a NEXT Configuration is built from a Configuration that has been activated.

When a NEXT Configuration is activated:

- The ACTIVE Configuration and Policies become the BACKOUT Configuration and Policies.
- The NEXT Configuration and Policies become the ACTIVE Configuration and Policies.
- DBS takes additional action to account for jobs that are awaiting execution:
  - Jobs that require devices that do not exist in the NEXT Configuration are detected and placed in MHS_TM HOLD for DBS.
  - Jobs that are already executing under DBS management are allowed to continue.

In both cases, DBS automatically activates the **BASE** Policy of the new ACTIVE Configuration. To activate a different Policy, follow the procedure described in Activating a Policy. Note that because you are activating a new Policy, you are also must choose between NORMAL and PASSIVE mode. See Activating a Policy for more details.

Creating, Editing, and Replacing the NEXT Configuration

The ACTIVE Configuration cannot be edited, therefore to make changes to a Configuration you must use the **CN** command to create the NEXT Configuration.

The **CN** command results in the Edit Configuration menu. *Note that if you exit from this menu without making any changes, you have not actually created the NEXT Configuration.*

Because DBS cannot anticipate how your Configuration changes affect existing Policies, the newly created NEXT Configuration contains only the default **BASE** Policy. *Any policies that were added to the ACTIVE Configuration, as well as any modifications that were made to the ACTIVE Configuration's **BASE** Policy, are not copied.*

Upon creating the NEXT Configuration, you are given the opportunity to modify the new **BASE** Policy if necessary. Again, we recommend that you do *not* change the **BASE** Policy without careful consideration. This ensures a consistent starting point to develop your specialized Policies, as well as an easily restored and predictable Policy. Instead, we recommend that you use the **MN** command, which allows you to add and edit Policies. See Managing Policies below for further details.

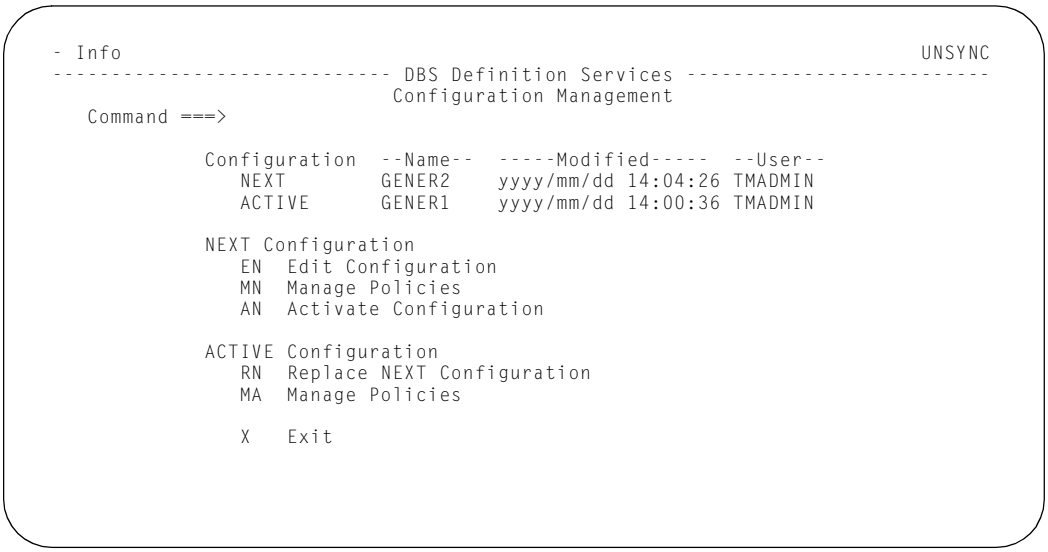

Until you activate the NEXT Configuration, you can use the **EN** command to edit it. If you want to start over from the beginning, you can use the **RN** command, which will completely replace the NEXT Configuration with a copy of the ACTIVE Configuration. Just as with the **CN** command, the only Policy associated with the replacement NEXT Configuration is the default **BASE** Policy.

Backing Out of a Configuration

The BACKOUT Configuration becomes available once you have activated the NEXT Configuration. The **AB** command restores the previous ACTIVE Configuration and Policies. If the NEXT Configuration exists, it is deleted, since it is based on the ACTIVE Configuration that is being abandoned. DBS automatically activates the **BASE** Policy of the new ACTIVE Configuration. Again, because you are activating a new Policy, you are given the choice of using NORMAL or PASSIVE mode. See Activating a Policy below for more details.

Managing Policies

Policies are associated with a Configuration and are more flexible in terms of management. For each Configuration a **BASE** Policy is created. Any number of Policies (well, almost any number) can be created and named. They are always constructed under the framework of a Configuration and are permanently associated with it.

Creating Specialized Policies

To create a specialized Policy, begin at the Configuration Management menu. Select **MN** to manage your Policies. This takes you to the Policy Management screen, which displays a list of your

installation's Policies. When you have just completed defining your first Configuration, the only Policy shown here is **BASE**.

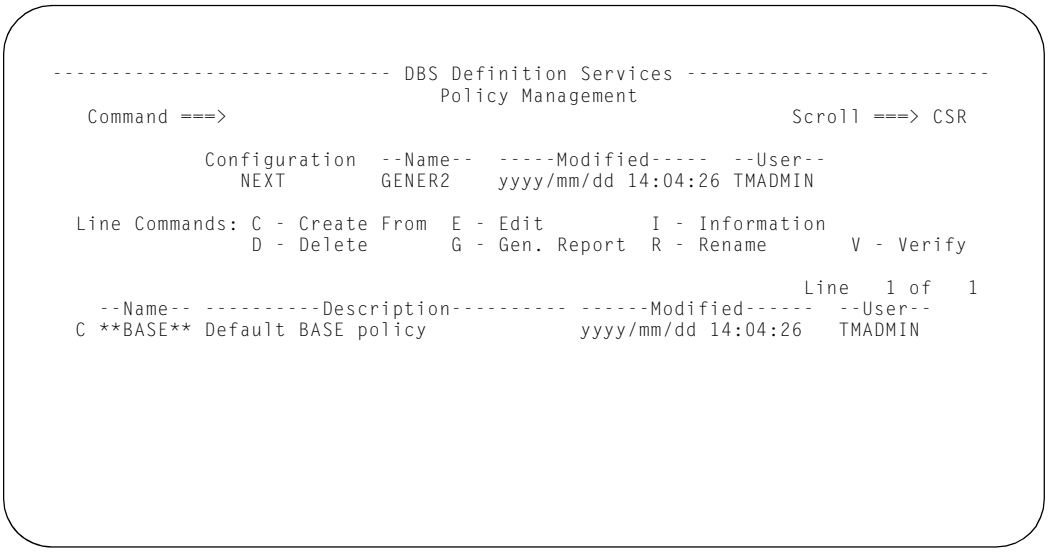

The **C** line command creates a new Policy based on the selected Policy. You are prompted to provide a new Policy name and description. You are next given the opportunity to manage your device counts so that they reflect your specialized Policy. After copying a Policy, the Policy Edit menu allows you to modify the device counts and turn the Work Groups feature on or off. See Configuring Work Groups.

Setting Device Counts

The default **BASE** Policy assumes that:

- All device counts are equal to the number of Device Numbers that you specified in your Configuration.
- All devices are shared as described in your Configuration.

To change the device counts available to a particular JESplex member, select the Manage Device Counts entry on the Policy Edit menu and use the **E** line command for the Drive Pool(s) that you want to edit. In the resulting pop-up window, tab to or click on the JESplex member to be changed.

Another pop-up window opens, in which you must indicate whether you want to use the default count (**Y**) or a modified count (**N**). If you select **N**, you are prompted to enter the new count. .

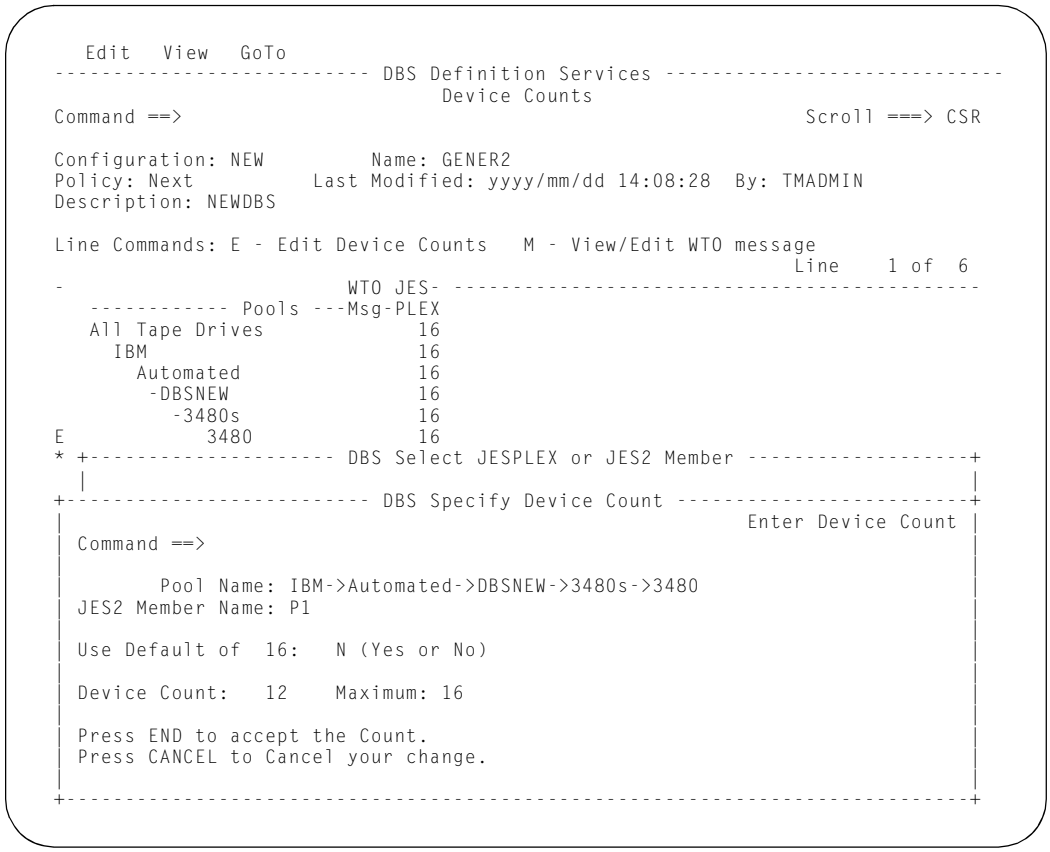

WTO Messages

For each Drive Pool, you can use the M line command to specify a WTO message that will be issued if the actual device count drops below the defined device count. This is detected during Policy activation or a DBS restart.

Activating a Policy

Policies can be activated at any time. When a new Policy is activated, DBS re-evaluates jobs that are awaiting execution and applies the new Policy to them. As a result, some jobs might now exceed the resources available to them under the new Policy. Use the Display of Jobs Exceeding Policy Resources in DBS Display Services, described in DBS Monitor and Displays to check for these jobs.

Running jobs that do not conform to the new Policy are permitted to proceed to completion.

PASSIVE Mode

When you ACTIVATE a Policy, you are asked to select Normal (default) or Passive mode. Passive mode means that DBS continues to evaluate jobs using the Configuration and Policy that was activated, but will not prevent any job or step from beginning execu- tion. Monitors, displays, and reports still reflect activity as though DBS were operating in normal mode, providing a means of evaluation and testing.

Configuring Work Groups

If you decide that Work Groups would be of benefit, you can turn the Work Group feature ON from the **Edit** menu on the Action Bar.

You can also use this menu to turn the Work Group feature OFF.

If you choose to turn the Work Feature OFF, you will lose all Work Group counts that have been previously defined.

Renaming Work Groups

Turning ON the Work Groups feature activates the ability to change the default Work Group names by using the **Edit** menu on the Action Bar. Select **Work Group Names/ Default/Options**, which displays the Work Group structure. Simply tab to a Work Group and type in a new name. The name must be alphanumeric, begin with an alphabetic or national ($, #, @) character, cannot exceed eight characters, and must be unique within the Policy. Note that these are the names that are specified when selecting a Work Group using JAL or JECL.

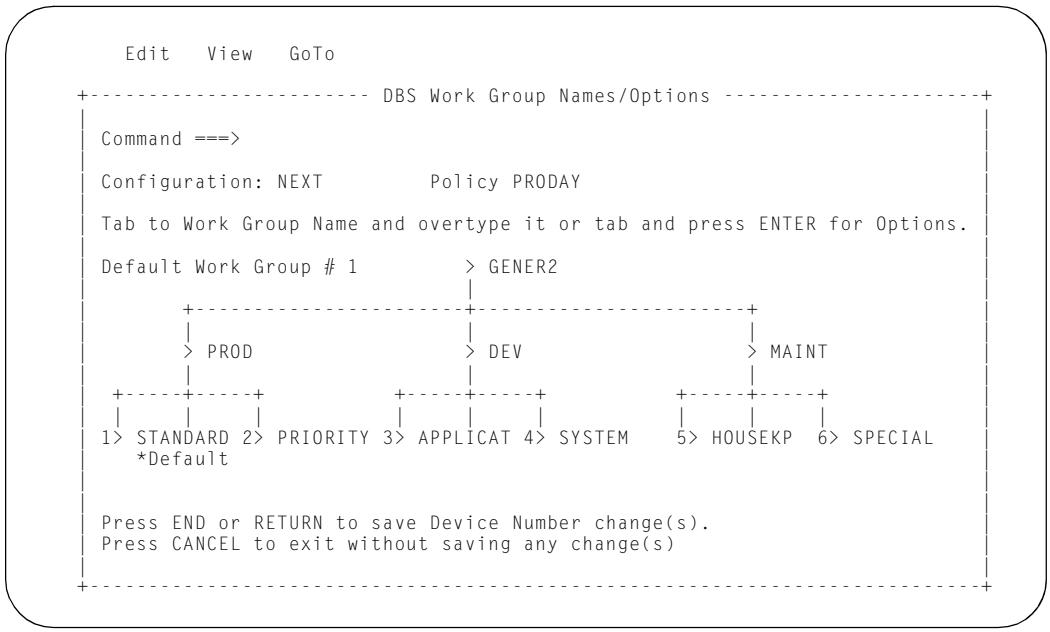

Setting the Default Work Group

You can change the default Work Group that DBS assigns if no Work Group is assigned in JAL or through JECL by changing the number specified in the **Default Work Group #** field.

Turning Overbooking Off

Overbooking allows DBS to make more efficient use of the devices in a Work Group at the expense of occasional allocation contention. If you have a Work Group for which allocation contention cannot be tolerated, you can turn off the overbooking algorithm by tabbing to the Work Group and pressing **Enter**. This opens a pop-up window that allows you to control overbooking for that Work Group.

Setting Work Group Counts

Before beginning to set Work Group counts, we strongly recommend that you study carefully the sections Work Group Planning and The Definition Process: What Should I Have? Remember that the dialog is intended for data entry, not planning.

The dialog has been designed to eliminate confusion as much as possible, therefore a flow is enforced. For example, you cannot dedicate devices to a Work Group for a JES2 member without first dedicating drives to the Work Group at the JESplex level. Similarly, you cannot dedicate drives to a Level 2 Work Group unless drives have already been dedicated to the Work Group's Level 1 parent.

Using Installation Pools

Installation Pools provide flexibility. All of your tape devices can be managed using the Drive Pools, but your installation might have special requirements. For example, if you have an esoteric unit name

that includes different types of tape devices, you could manage them as a group by using an Installation Pool.

To define an Installation Pool, you must first name the library you will use. Installation Pools support two types of library:

- Esoteric, which you can use to map any esoteric unit names your installation might use.
- Non-esoteric, which you can use to accommodate special circumstances that cannot be handled any other way.

After you have named the library, enter the device numbers in the usual way. Note the following:

> Any device number added to an Installation Pool must already exist in a Drive Pool. This is an exception to the rule that a device number must be defined in one Pool only.

For information about assigning Installation Pools using DAL, see Assigning Installation Pools.

Using JAL and JECL with DBS

DBS does not require DAL or JAL to deliver its basic benefits; however, you will probably need DAL/JAL if you want to take advantage of Work Groups.

Assigning Work Groups

Although Work Groups can be assigned using JECL, they are usually assigned in JAL with the DBS SET statement:

```
DBS SET WORKGROUP(PROD)
```

The name that you specify must be one of the six Subgroups specified in the DBS Policy.

Setting DBS Priorities

In an environment without SLM, DBS Priorities are also usually assigned using the JAL statement DBS SET:

```
DBS SET PRIORITY(HIGH)
```

There are three DBS Priorities:

- LOW
- MEDIUM, which is the default
- HIGH

The scope of DBS Priorities is the JESplex.

These Priorities do not affect job selection. Rather, they determine which job gets drives when contention occurs. This decision is made during step allocation.

Note that if your installation is using SLM, it will manage priorities rather than DBS.

Implied Priority

If no DBS Priority is assigned or if jobs have equal Priorities, preference is determined by the assigned Work Group. Lower numbered Work Groups, as shown in a display of the DBS Work Groups structure, are served first.

Assigning Installation Pools

Installation Drive Pools are used to define pools of tape devices that your installation has grouped at some level other than their generic device type, such as those for which your installation has created

esoteric unit names. DBS manages Installation Pools in the same way that it manages normal Drive Pools. Installation Pools are assigned using the DAL statement DBS ASSIGN:

```
DBS ASSIGN INSTALLATION_POOL(TAPE)
```

The DBS ASSIGN statement is the only way you can assign tape devices to Installation Drive Pools.

Unavailable Unit Handling

The way DBS handles jobs that request more devices than have been defined in the Configuration can be changed with the JAL statement DBS HOLD:

```
DBS HOLD UNAVAILABLE_UNITS(YES)
```

Jobs that request more units than have been defined are normally failed with a JCL error. This statement requests that they be placed in an MHS_TM HOLD instead.

Note that such jobs can never run under the ACTIVE Configuration.

JECL for STCs

The JECL statements for DBS are intended primarily for STCs, which do not pass through Job Analysis. Although you can use these statements in ordinary batch jobs, the information provided by them should normally be provided using DAL/JAL, which always supercedes JECL.

Reserving Devices With //*+DBS RESERVE

Jobs that use tape devices without being registered with DBS are considered poachers. Because DBS has no information about poachers until they actually allocate devices, it is harder to account for their tape device usage. A job is potentially a poacher when:

- The job has one or more steps that use dynamic allocation.
- The job is a started task, which by its nature does not undergo Job Analysis.
- The job is made exempt from Job Analysis, either by its submission class or by an operator command.

You can avoid poaching and register STCs and jobs with DBS by including the //*+DBS RESERVE JECL statement:

```
//*+DBS RESERVE IBM->AUTOMATED->ATLDS1->3480S->3480=(1,1)
```

This example is for an STC, and informs DBS that the STC will request one device through dynamic allocation (the first number) and one through static allocation (the second number). For batch jobs that pass through Job Analysis, the second number is ignored, since ThruPut Manager calculates the number of static allocations.

Setting Work Groups and Priorities With //*+DBS SET

You can use the JECL statement //*+DBS SET to assign Priorities and Work Groups:

```
//*+DBS SET WORKGROUP=SPECIAL,PRIORITY=HIGH
```

Accounting for Dynamic Allocation

ThruPut Manager determines resource counts by analyzing the job's JCL, but this method cannot account for dynamically allocated resources. In order to maintain accurate resource counts, use the DAL statement DBS RESERVE:

```
DBS RESERVE IBM->AUTOMATED->ATLDS1->3490->3490E(1)
```

This statement adjusts the resource count for the IBM Automated library ATLDS1 3490E Drive Pool. Note that this statement also increases the resource counts of JAL Descriptors describing the specified unit type.

DBS Monitor and Displays

DBS provides a rich selection of monitor and display screens, as well as a report for a particular Configuration/Policy combination. There are far too many variations to include examples of each one, but the following should give you an idea of the power available from DBS Display/Monitor Services.

To invoke DBS Display/Monitor Services, first issue the TSO command under ISPF to invoke ThruPut Manager ISPF services:

```
TMISPF
```

This takes you to the Main Lobby for ThruPut Manager dialogs. Select, in order:

1. Display/Monitor Services
2. Automation File Services
3. Exit

You can now choose DBS - Drive Booking Services Display and then see options of DBS Display Services or DBS Monitor Services.

DBS Display Services

Invoking DBS Display/Monitor Services collects a snapshot of data relevant to DBS Displays. All DBS Displays show the timestamp reflecting when this snapshot was taken. To refresh the data, use the REFRESH primary command, or select the **Refresh** item from the **View** drop-down menu on the Action bar.

Using DBS Display Services, you can:

- Display the details of the ACTIVE Configuration and Policy.
- List any jobs that are requesting more devices than are available in the current Policy.

To explore the power of DBS Display Services, we recommend that you experiment with selecting JES2 members, Work Groups, and with the various line commands. The online help (**F1**) explains each screen and its associated line commands.

DBS Monitor Services

Invoking DBS Monitor Services begins the data collection process for the various monitor displays. In other words, it is Monitor Services that collects the data for display, not the DBS application. The collected data is updated approximately once per minute.

To use DBS Monitor Services effectively, you must allow it to run for a while, then press **Enter** to refresh the screen. Accumulated values will then be displayed. Using DBS Monitor Services, you can monitor drive utilization by:

DBS Configuration/Policy Report

You can generate a DBS Configuration/Policy report from the Policy Management screen by using the **G** line command. This report can be useful for planning and documentation purposes. Excerpts from a sample report are shown on the following pages.

```
Drive Booking Services              DBS Configuration / Policy Report                          MON dd, yyyy  11:51
Page   1

                                        Configuration: CONFIG6      Policy: PRODDAY

Automation File: TM.AUTOFILE  On Volume: MVS003

GENERAL INFORMATION

Name: CONFIG6   Description: Sample of a Production Configuration

Created: yyyy/mm/dd  At: 11:06:31 Last Modified: yyyy/mm/dd  At: 11:47:57  By: TMADMIN

JES2 Node Name: TORONTO  Spool Dataset: SYS1.HASPACE  Volume Prefix: SPOO

JES2 Member Name(s): PR    T1    T2    T3    T4    T5    T6    T7    T8    T9    T10   T11   T12   T13   T14   T15

Vendor Options   EMC CopyCross Esoteric Unit Names:

                      ES031

                 Generic Unit Names:

                      GEN31

                 Subsystem Names:

                      SS3

   CA:VTape DATACLAS Names:

                      MEDIA1   MEDIA2   MEDIA3   MEDIA4

            INCLUDE_ON_VL(AUTOMATED)
```

DBS Configuration/Policy Report: General Information

```
Drive Booking Services              DBS Configuration / Policy Report                MON dd, yyyy  11:51
Page   2

                      Configuration: CONFIG6      Policy: PRODDAY

POOL SUMMARY


                                                 JES2 Member
 Pool                      Device Count  Override(s)

All Tape Drives               677
 IBM                           293
   Automated                   35
     ATLDS1                    35
       3480s                    3
         .3480                   2*          1
         .3480X                  1*          2
       3490                     32
         .3490E                  32*         2

  .....


* Counts entered in dialog.  All other counts are calculated.


Total Number of Pools: 37
 Total Number of Defined Pools: 14
 Total Number of Defined Pools with Member Overrides: 3

Total Number of Retired/Not-Used Pools: 5
```

DBS Configuration/Policy Report: Excerpt from the Pool Summary

Drive Booking Services DBS Configuration / Policy Report MON dd, yyyy 11:51
Page 3

Configuration: CONFIG6 Policy: PRODDAY

POOL DETAILS

Pool	Member	Device Count	Device Numbers
IBM-Automated-ATLDS1-3480s-3480	JESPLEX	2*	1580 159F
	T1	1*	159F
IBM-Automated-ATLDS1-3480s-3480X	JESPLEX	1*	1680
	PR	1*	1680
IBM-Automated-ATLDS1-3490-3490E	JESPLEX	32*	1780-179F
	PR	16*	1780-178F
	T1	16*	1790-179F

....

.....* Counts entered in dialog. All other counts are calculated.

DBS Configuration/Policy Report: Excerpt from the Pool Summary

Drive Booking Services DBS Configuration / Policy Report MON dd, yyyy 11:51

Page 4

Configuration: CONFIG6 Policy: PRODDAY

WORK GROUP SUMMARY

Default: #1 STANDARD

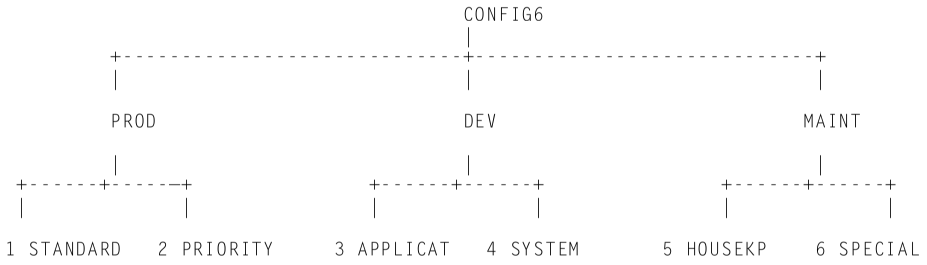

DBS Configuration/Policy Report: Work Group Summary

Appendix A. Supported Devices

This appendix provides a complete list of the devices supported by DBS, arranged by device type within mode within vendor. DBS is intended to support all tape devices that are supported by z/OS. If your installation uses a tape device that is not listed here, please contact ThruPut Manager Technical Support.

DBS Supported Devices			
Vendor	**Mode**	**Device Type**	**Model**
IBM	Automated	3480	
		3480X	
		3490	3490E
		3590-1	3590B
			3590E
			3590H
			3592J
	Manual	3400-N	
		3480	
		3480X	
		3490	3490E
			3592J
		3590-1	3590B
			3590E
			3590H
IBM	Manual	3590-1	3592J
	Virtual	3490V	3490E Virtual
StorageTek	Automated	3480	4480
		3490	4490
			9490
			9490EE
			SD3
			9840
			T9840B
			T9840C
			T9940A
		3590-1	SD3
			984035
			T9840B35
			T9840C35

DBS Supported Devices			
Vendor	**Mode**	**Device Type**	**Model**
			T9940A35
StorageTek	Non-library	3480	4480
		3490	4490
			9490
			9490EE
			SD3
			9840
			T9840B
			T9840C
			T9940A
		3590-1	SD3
			984035
			T9840B35
			T9840C35
StorageTek	Non		T9940B35
StorageTek	Virtual	3490V	3490E Virtual
CA-Vtape	Virtual	3480V Virtual	
		3490E Virtual	
EMC CopyCross	Virtual	3480V Virtual	
		3490E Virtual	