Compuware | The Mainframe Software Partner For The Next 50 Years

ThruPut Manager AE
Automated Capacity Management Guide

Release 18.02

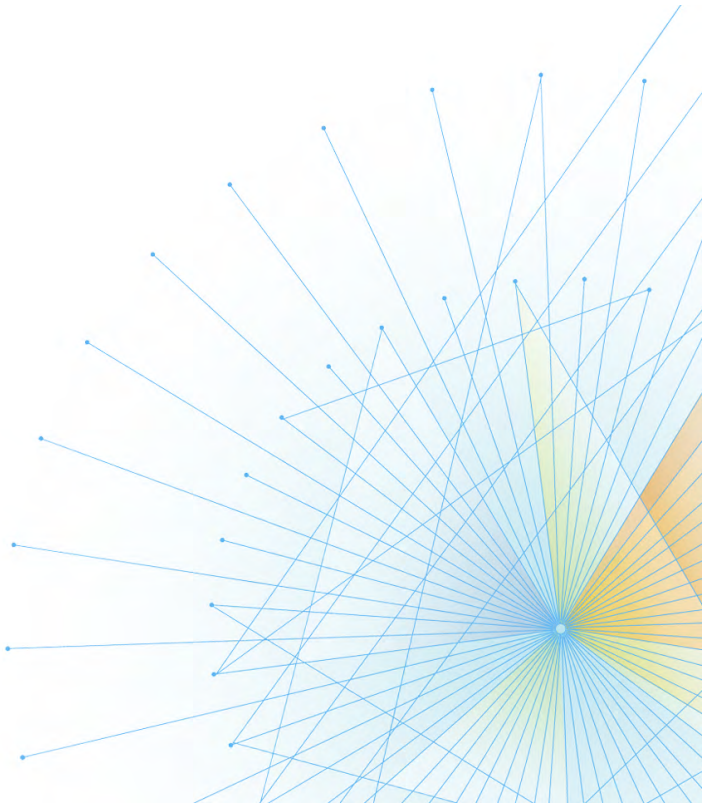

Please direct questions about ThruPut Manager
or comments on this document to:

ThruPut Manager **Customer Support**

https://go.compuware.com/

Doc. APR2019
March 1, 2019

Introduction

Summary of Changes

V1802-7118 (*April 2019*)	• ACM now supports discounting of the R4HA of mobile and public cloud workloads when classified appropriately by WLM. 60% of the R4HA for mobile and public cloud is subtracted from the total R4HA for an LPAR Set before comparing the total against the installation supplied limit.
V1802-7117 (*January 2019*)	• ACM now supports the Estimated Time to Capping values calculated by WLM for z/OS 2.3 and above.
V1802-7116 (*October 2018*)	• Country Multiplex Pricing (CMP) - Multiplex Scope. CMP LPAR Sets usage may now be accumulated across JESplexes on all sysplexes.
V1802-7115 (*July 2018*)	• Country Multiplex Pricing (CMP) - Sysplex Scope. CMP LPAR Sets usage is now accumulated across JESplexes within the same sysplex.
V1802-7114 (*April 2018*)	• No changes
V1802-7113 (*January 2018*)	• Support for Country Multiplex Pricing (CMP). LPAR Sets may now be defined as CMP LPAR Sets and have a CMP Limit for the R4HA.
V1802-7112 (*October 2017*)	• Integrated Automated Capacity Management for Production Services workload into this guide.
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
V7R1-7108 *(February 2017)*	• The topic Update Your ThruPut Manager AE Policy For ACM was edited to reflect the presence of Quick Jobs and Production Services Groups.
V7R1-7107 *(May 2016)*	• No changes
V7R1-7106 *(November 2015)*	• The definition and implementation of LPAR Sets were added. – LPAR Sets – Implementing LPAR Sets • The description of how ACM coexists with 3rd party capping products was added. – TM AE and 3rd Party Capping Products – Running ACM with – 3rd party capping software • The following topics were edited to reflect the presence of LPAR Sets and 3rd party capping products. – Reducing R4HA without Soft Capping – Update your ThruPut Manager AE policy for ACM

V7R1-7104 *(July 2015)*	• The example of setting up ACM constrains (Chapter 2) was restructured
V7R1-7101 *(July 2014)*	• This is a base manual for ThruPut Manager Version 7 Release 1.0.

About This Manual

Automated Capacity Management (ACM) is an elective feature of ThruPut Manager AE that is used in conjunction with IBM's sub-capacity pricing to reduce software costs. It monitors the rolling 4-hour average (R4HA) of CPU usage, the basis for software charges, and controls batch workload in order to reduce that average. This can significantly reduce software MLC charges. Also, by reducing workload before the limits are reached, you can avoid "running into the wall'---hitting the limit with a very large concurrent workload which can cause soft capping actions to be severe enough to cause problems with online response.

When used together with soft capping, ACM minimizes the effects of reaching the capping limit or avoids capping altogether, and enables your datacenter to lower your R4HA limit without impacting your online and critical batch systems.

This manual introduces the underlying concepts of the ACM feature of ThruPut Manager AE, guides you through the setup process, and describes runtime interaction. It assumes you have completed an "essential implementation" and, therefore, have created a Policy with Service Groups to describe and control the processing of workload subject to Automated Capacity Management. This guide is of interest to the ThruPut Manager technical support person, who acts in consultation with performance and capacity management specialists.

This manual assumes that you are familiar with the contents and examples contained in the *ThruPut Manager AE Concepts Guide* and the *Usage Guide*.

Contents

Chapter 1
ACM Concepts

ThruPut Manager AE monitors the rolling 4-hour average and defers batch workload when appropriate, resulting in reduced charges in a sub-capacity pricing environment.

Introduction

Automated Capacity Management (ACM) is an elective feature of ThruPut Manager AE that is used in conjunction with IBM's sub-capacity pricing to reduce software costs. It monitors the rolling 4-hour average (R4HA) of CPU usage, the basis for software charges, and controls batch workload in order to reduce that average. ACM builds on the TM AE Policy by specifying constraints that are applied as the R4HA approaches the capacity limit.

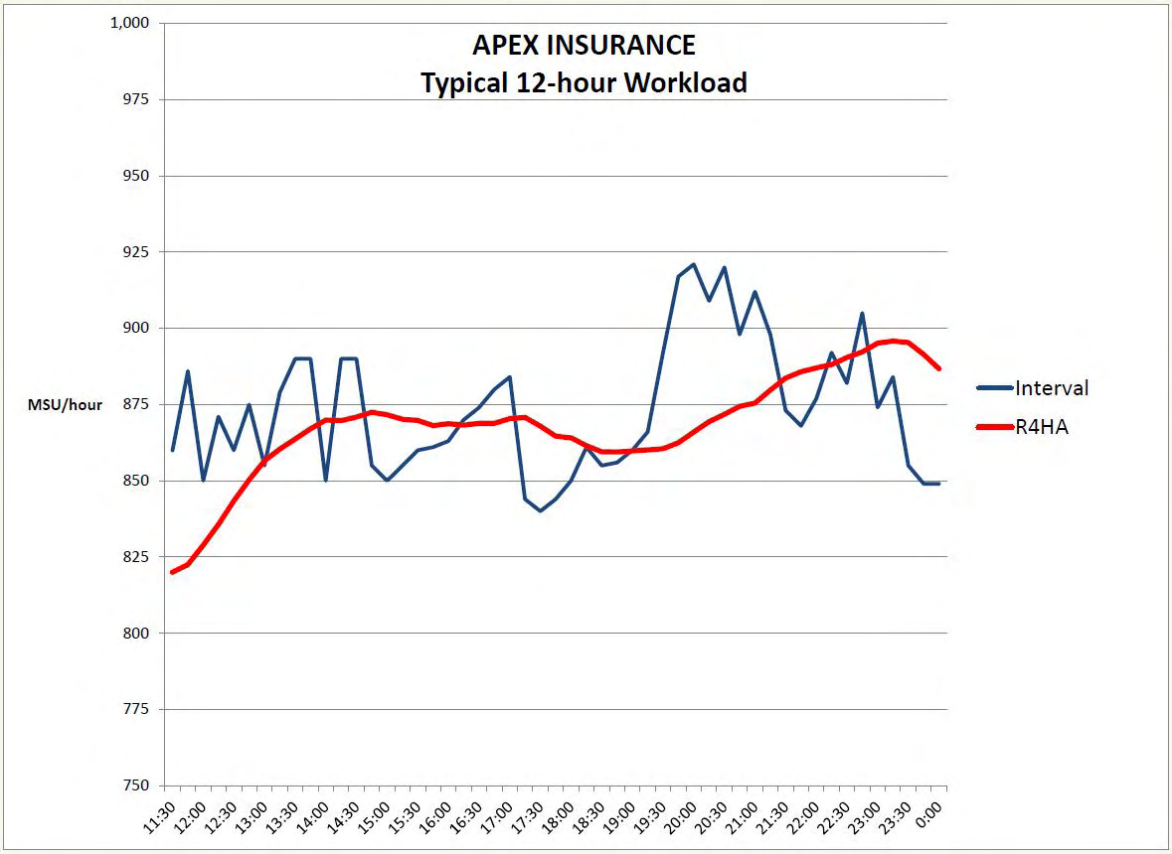

This chart shows that, while the interval CPU usage may fluctuate significantly, the R4HA changes slowly. The peaks in the interval usage do not necessarily coincide with the peaks in the R4HA. Sub-capacity pricing is based on the highest monthly peak in the R4HA, not the interval peak.

When used together with soft capping, ACM minimizes the effects of reaching the capacity limit or avoids capping altogether, and enables your datacenter to lower your R4HA limit without impacting your online and critical batch systems.

The scope of ACM functionality is the LPAR. Since the scope of a ThruPut Manager AE Policy is a JESplex, ACM settings within the Policy are applicable to each LPAR within the JESplex.

Though not mandatory, results are optimized when ACM is enabled on all the LPARs within a CPC.

Batch Processing And Sub-Capacity Pricing

When your datacenter has sub-capacity pricing agreements, your software licensing charges for a month are based on that month's highest rolling 4-hour average (R4HA) CPU usage, the R4HA peak. (Peaks in the (interval) CPU utilization are not a good predictor of the R4HA peaks.)

Your datacenter can control the R4HA peak by setting, in the HMC, a LPAR Group Limit for a group of LPARS or Defined Capacity limit for a particular LPAR, or both, giving an effective capacity limit. When the R4HA reaches the capacity limit, service given to the running workload is reduced as required until the R4HA falls below the limit. This is called soft capping the workload.

The drawback to soft capping is that it does not differentiate between types of workload. All workload demanding CPU cycles is restricted, even online and critical batch. Therefore, datacenters keep the capacity limit high enough to protect their online workloads.

LPAR Sets

Some installations have requirements for which Defined Capacity or LPAR Group Limits are not sufficient. For example:

- Define a limit on a group of LPARs without implementing soft capping.
- Establish separate limits for various subsets of LPARs based on the differing sets of software products that they run. Some LPARs might belong to more than one subset.
- One or more LPARs must not be capped but their usage should be included in R4HA calculations. Some LPARs may be hard capped and therefore cannot be subject to Defined Capacity or LPAR Group Limits.
- The installation makes use of Country Multiplex Pricing (CMP) and wants to control batch workload based on an overall limit for the R4HA across multiple CPCs on any or all JESplexes in all of the Sysplexes in the multiplex.
- The installation takes advantage of Mobile Workload Pricing or has public cloud workloads.

ACM enables the installation to define subsets of the z/OS LPARs on each CPC as named LPAR Sets and assign a R4HA limit to each LPAR Set. On each LPAR on which ACM is active, ACM monitors and calculates the R4HA for each defined LPAR Set to which the LPAR belongs.

LPAR Sets can be used whether or not the installation has specified Defined Capacity or LPAR Group Limits. ACM will automatically determine all the applicable LPAR Set limits for each LPAR and factor these into its calculations. Note that an LPAR Set limit being reached does not cause soft capping to occur.

An LPAR Set may also be defined as a CMP LPAR Set. When specified, ACM will total the R4HA on every CPC in all JESplexes in the Sysplex on which the LPAR Set is defined and compare that total usage against a specified overall CMP limit.

Note that a CMP LPAR Set must be defined with the same name in all JESplexes in all Sysplexes.

In order for ACM to correctly identify each instance of the CMP LPAR Set, it must be defined with the same name in each active SLM Policy on each JESplex in the multiplex.

If you are uncertain as to what value to use for the limit for a given LPAR Set, you can specify zero. While ACM will monitor the LPAR Set and calculate the R4HA, it will take no actions based on the R4HA for the LPAR Set while the limit remains zero. TM AE provides operator commands as well as SMF records that allow you to monitor the R4HA for all LPAR Sets, providing information that will

prove useful in setting a limit value. This applies to both regular and CMP LPAR Set limits. You may wish to always leave the regular LPAR Set limit at zero and set only a CMP limit.

When defining a CMP LPAR Set in multiple SLM Policies across multiple JESplexes, you must ensure that for each CPC, the same LPARs are defined as belonging to the LPAR Set and that the overall CMP limit for the LPAR Set is the same. If there are differences between the definitions for one or more JESplexes, ACM will issue messages every 15 minutes that indicate either a difference in LPARs defined for the LPAR Set on an indicated CPC, or that the limits are inconsistent.

If there is a conflict in the limit specified for a given LPAR Set, ACM will use the limit from the most recently updated SLM Policy. This is designed to accommodate the adding a new CPC to a sysplex. When this occurs, you may wish to increase the limit for the affected LPAR Sets. Just set the new limit in the SLM Policies for the JESplexes that are on the new CPC, and ACM on the other JESplexes will automatically use the new limit. They will issue warning messages every 15 minutes and when you have time you should update those policies as well.

Using CMP LPAR Sets across Sysplexes

Within a sysplex, ACM uses XCF to exchange LPAR Set usage data between JESplexes and therefore requires no additional configuration.

However, ACM uses TCP/IP to communicate between JESplexes that are on different sysplexes. Therefore, in addition to the normal z/OS TCP/IP configuration requirements, you will need to supply additional 'TM TCPIP' TMSS initialization statements.

You need to identify the port number to be used by the local ThruPut Manager to receive messages (TM TCPIP SET PORT) as well as provide at least one IP address and port number for every JESplex on another sysplex with which ACM will exchange information on LPAR Set usage (TM TCPIP ADD). It is strongly recommended that you supply this information for every member of every JESplex so that ACM can still communicate as long as one member in a JESplex is up.

See the *ThruPut Manager Base Product System Programming Guide* for details on the TM TCPIP SET and TM TCPIP ADD statements.

You can also use TM TCPIP commands to supply information on remote JESplexes dynamically until you have the opportunity to restart ThruPut Manager. See the *ThruPut Manager Command Reference Guide* for details.

Note: The port number for the local ThruPut Manager to use to receive messages can only be supplied using a TMSS initialization statement (TM TCPIP SET PORT). It cannot be set via command.

Mobile and Public Cloud Workload Pricing Discounts

ACM supports discounting the mobile and public cloud contribution to the R4HA of any CMP LPAR Set as specified in the following IBM pricing options: Mobile Workload Pricing (MWP) and Z Workload Pricing for Cloud (ZWPC).

You must specify in the SLM Policy that you want to activate discounting for mobile and/or public cloud.

When discounting is activated, for every CMP LPAR Set, ACM will automatically accumulate the R4HA and 5 minute interval consumption rate for all workload classified by WLM as “MOBILE” or “CATEGORYA” (public cloud), and then subtract 60% of the accumulation from the R4HA and 5 minute interval consumption rate totals for the LPAR Set.

This enables you to set the limits for each LPAR Set to your desired targets based on the discounted peak values that are reported by the SCRT. ACM will automatically consider the amount of mobile and pubic cloud workload when determining how much batch workload can be run given each LPAR

Set limit. This means that the same limit applies all the time, regardless of how much mobile or public cloud workload may be running.

Note that, in order to obtain complete information on mobile and public cloud workload consumption, ThruPut Manager with ACM enabled must be running on all z/OS LPARs.

This support is only available for CMP LPAR Sets as it relies on information being exchanged between LPARs that may not be in the same JESplex or Sysplex. ACM only exchanges this information for CMP LPAR Sets. ACM will add up the mobile and public cloud workload usage rates across all LPARs on all CPCs.

However, it is not necessary to use Country Multiplex Pricing to take advantage of support for discounting workloads. By giving each CMP LPAR Set a name that is unique to each CPC, the discounts and usage will be totaled across LPARS on the same CPC only.

Constraining Batch Workload

A typical datacenter has a batch workload that accounts for 25%-40% of CPU usage, with online processing accounting for most of the rest. Of the batch workload, typically 25% is high importance, 50% is medium, and 25% is low importance.

ACM builds on the TM AE Policy information to identify batch workload that can be deferred as the R4HA approaches its capacity limit. You choose how to defer this workload by specifying ACM-specific constraints for the relevant Production or General Service Groups or, in the case of Production Services workload, for the relevant Production Importance.

At run time, ThruPut Manager AE monitors the R4HA for the current LPAR and any LPAR Group and LPAR Sets to which it belongs. ACM constrains the batch work you identified based on the usage of the LPAR Group, LPAR Set or LPAR that ACM determines is the highest compared to its limit.

ACM Acts Before the R4HA Reaches Its Limit

In a series of steps, Automated Capacity Management puts progressively more constraints on which jobs are selected, and, once executing, on the amount of CPU they are allowed to consume. These steps, or Capacity Levels, are expressed as a percentage of the limit.

Capacity Level	% of Soft Cap (default)
1	99%
2	95%
3	90%
4	85%
5	80%

Jobs with constraints associated with Capacity Level 5 have these constraints applied when the R4HA reaches 80% of the limit. The Capacity Level 4 constraints are applied when the R4HA reaches 85%, and so on. Correspondingly, as the R4HA decreases, ACM relaxes the constraints, level by level, until, below Capacity Level 5, it imposes no ACM constraints at all, and all batch jobs run normally.

You can modify the percentages for Capacity Levels 2 to 5 should the default values prove inappropriate for your datacenter.

Batch jobs typically use considerable amounts of CPU, so constraining them in progressive steps keeps your overall CPU usage under control, without triggering capping.

ACM Supports Estimated Time to Capping

On z/OS 2.3 and higher systems, ACM supports the estimated time to capping values calculated by WLM. These are WLM's approximation of when soft capping will begin either for the LPAR or the LPAR Group.

When the estimated time is less than the lead time value supplied in the RTCapLeadTime parameter in the active IEAOPTxx member of the z/OS PARMLIB, if enabled in the active SLM Policy, ACM will force the Capacity Level value to 1. This will activate the maximum constraints on batch workload specified by the installation in the active SLM Policy.

Note that the IBM-supplied default value for the RTCapLeadTime parameter is 0 so, by default, ACM will ignore the estimates. However, if you do not want to set RTCapLeadTime or want to use a different value for ACM's use, you can specify a lead time value in the SLM Policy, which ACM will use instead of the value of RTCapLeadTime.

ACM Constrains Subsets of Batch

ACM reduces service to subsets of the batch workload as each Capacity Level is reached. If a job is not subject to ACM constraints as it would without ACM enabled.

> Note: All Workload Adds to the R4HA. No matter how unimportant a workload is, the CPU cycles it consumes contribute equally to the R4HA.

There are three ways you can reduce service to a job when a Capacity Level is reached.

Reduced Service Class

When a capacity level is reached, jobs use ACM-specific Service Classes with Resource Group maximums so that the jobs consume limited resources.

This constraint can be applied to all PS work with a specific importance, jobs assigned to a PS Group, jobs assigned to a set of GS Groups (identified by a Control Center), or jobs assigned to a specific GS Group.

Note when a job has production importance for which a reduced service class is specified, and belongs to a PS group for which a different reduced service class is specified, the one from the PS Group is used.

Job Selection Point

This option allows the installation to specify the earliest threshold that a job must reach in order to be selected once a Capacity Level has been reached. (Recall that a threshold is a point in the queue described by the target-acceptable-critical scale.) In effect, jobs stay in the queue longer before being selected.

This constraint can be applied to all PS work with a specific importance, jobs assigned to a set of GS Groups (identified by a Control Center), or jobs assigned to a specific GS Group.

Maximum Number of Jobs

There is a maximum number of jobs that can be active (executing) at a time, on a JESplex member once a Capacity Level was been reached. If the maximum is reached TM AE does not select new jobs for this member until some of the active jobs finish.

This constraint can be applied to jobs assigned to a PS Group, jobs assigned to a set of GS Groups (identified by a Control Center), or jobs assigned to a specific GS Group.

Your ACM constraints should be stricter than those already specified in your TM AE Policy, since these ACM constraints are applied in addition to and not instead of them.

Reducing R4HA without Soft Capping

If your datacenter does not wish to enable soft capping, you can still benefit from Automated Capacity Management.

By defining LPAR Sets, each with their own limit and/or CMP limit, you provide ACM with the information needed to calculate the Capacity Level and constrain batch workload based on the active TM AE Policy.

If you wish to establish a R4HA limit for each LPAR, you can supply one in the TM AE Policy using the dialog. Note that these limits for each LPAR are ignored if either Defined Capacity or an LPAR Group Limit is defined using the HMC.

So even without either a Defined Capacity or LPAR Group Limit being specified, ACM still calculates the Capacity Levels using LPAR Set limits and any LPAR limit that has been supplied in the TM AE Policy. ACM then will constrain batch work using the constraints you have specified for the Capacity Level.

Reducing MLC Charges

After some experience with Automated Capacity Management, you may consider reducing your Defined Capacity, LPAR Group Limits and/or LPAR Set Limits, and so reduce your monthly software licensing charges, while ensuring your online and high importance batch work is not affected adversely.

For further details regarding the technical aspects of ThruPut Manager processing in a sub-capacity pricing environment, see the *Automated Capacity Management Technical Background* white paper, available upon request.

TM AE and Third Party Capping Products

You may already have a third party capping product installed that is continually adjusting the Defined Capacity limit for each LPAR and/or the LPAR Group Limit.

In order for ACM to work with these third party capping products, you provide extra information in the TM AE Policy.

If a third party capping product is changing the Defined Capacity and/or LPAR Group limit frequently, ACM can no longer use the limit(s) for comparison. Instead, using the TM AE dialog, you set options in the policy to tell ACM to ignore the limit(s) being adjusted by the third party product.

ACM will then use LPAR Set limits and individual LPAR limits set in the policy to determine the Capacity Level, allowing ACM to continue to reduce the R4HA by constraining the batch workload as you have specified in the policy. If the third party capping product only changes the Defined Capacity, then ACM can also continue to track the usage for an LPAR Group if defined.

Contact Compuware Customer Solutions for assistance in configuring ACM to coexist with a third party capping product.

Chapter 2
Set Up ACM Constraints

You set up Automated Capacity Management by designing your ACM constraints, specifying low-resource-consumption WLM Service Classes for ACM use, and specifying your constraints through the ThruPut Manager AE dialogs.

Design Your ACM Constraints

As discussed earlier, Automated Capacity Management has five Capacity Levels, each defined as a percentage of the capacity limit. You identify the workload to which you are prepared to give reduced service as each of these Capacity Levels is reached and then decide which ACM constraints to apply.

You can specify a combination of constraints for a single Service Group, to both slow down selection and give reduced service to jobs when they are selected.

Begin by looking at your Service Group definitions. Importance and service time goals can help guide this analysis. Consider what work should be constrained at Capacity Level 5 and how it should be constrained, then continue for Capacity Level 4, and so on.

You'll find it useful to record your decisions as you go through the process of designing your ACM constraints, bearing in mind some work is considered business-critical and so no constraints are specified.

If some work, assigned to a Service Group with ACM constraints, should not be constrained even if the rest of the jobs belonging to this Service Group are, then it can be identified in JAL rules and excluded from Automated Capacity Management constraints by specifying SLM SET NO_CAPACITY_MANAGEMENT.

Specify ACM Service Classes in Your WLM Policy

With Automated Capacity Management enabled, you have the opportunity to direct a job to a different set of WLM Service Classes when your system has reached any of its Capacity Levels, e.g., TMGSCL1, TMGSCL2, TMGSCL3, TMGSCL4, and TMGSCL5 for jobs in the General Services queue and TMPSCL1, TMPSCL2, TMPSCL3, TMPSCL4, and TMPSCL5 for jobs in the Production Services.

> Note: Only one of these Service Classes is in use at any one time, so WLM has only at most two additional active Service Class period to manage.

Design these ACM-specific Service Classes so that they have a lower resource consumption rate than the ones assigned for normal ThruPut Manager AE operation. Jobs running in these Service Classes have limited access to runtime resources, thus helping to lower CPU usage.

Plan for one Service Class for each Capacity Level, with the following characteristics:

- Each should be discretionary.
- Each should belong to a WLM Resource Group with a defined maximum at the system level.
- Resource Group maximums should decrease at each step, as the associated Capacity Level goes from 5 to 1.

For example, ensure the maximum for the Resource Group for the Service Class to be used at Capacity Level 2 is lower (more restrictive) than the maximum the Resource Group for the Service Class to be used at Capacity Level 3.

Make ACM Updates in Your TM AE Policy

Once you've designed your Automated Capacity Management constraints and modified your WLM policy to include ACM-specific Service Classes you need to:

1. Create a copy of your existing ThruPut Manager AE policy to edit.
2. If you wish to use LPAR Sets, define the following:
   a. Each of your CPCs.
   b. The LPARs running z/OS on each CPC.
   c. LPAR Sets by selecting the LPARs for each set, and assigning a name and a limit. You may also designate the LPAR Set as a CMP LPAR Set and specify a CMP limit.
3. If you have a third party capping product active, set the appropriate options in the TM AE Policy. Contact Compuware Customer Solutions for assistance.
4. Map the ACM-specific Service Classes to the Capacity Levels in your TM AE Policy.
5. Specify ACM constraints for jobs with a specific Production Importance.
6. Specify ACM constraints for jobs assigned to a Production Service Group.
7. Specify ACM constraints for jobs assigned to a General Service Group.
8. Enable ACM constraints.
9. Decide if Quick Jobs should be constrained along with your other work.
10. If your datacenter does not use soft capping, LPAR Sets are usually sufficient to control the batch work based on R4HA. However, if you do require a limit per LPAR, specify a capacity limit per LPAR in the TM AE Policy.
11. Optionally, adjust your percentages for Capacity Levels 2 to 5. This is only required if the defaults are inappropriate.
12. If necessary, add JAL to exempt specific SLM-managed jobs from ACM control.
13. Activate the updated policy.

1. Create a Copy of Your Policy

Use the Policy Management panel to make the copy of your policy. You'll update this copy and later activate it to put your changes into effect.

2. Implement LPAR Sets (Optional)

Implementing LPAR Sets requires the following three steps:

a. Define your CPCs.

b. For each CPC, define all LPARs that are running z/OS.

c. For each CPC, define the LPAR Sets that you wish to create, selecting from the LPARs that you have defined for the CPC. Supply a name and a limit for each LPAR Set. You may also designate it as a CMP LPAR Set and assign a CMP limit. If you are implementing ACM across multiple JESplexes, make sure that each CMP LPAR Set is defined as having the same LPARs on each CPC, and that the overall CMP limit is the same.

a. Define your CPCs

You need two pieces of information to define a CPC: the type and the serial number. You can find both of these on the SCRT report, but there's an easier way to do it. Just issue the TM AE command /SLM D CPC on an LPAR on the CPC you wish to define, and the command response will display the CPC type and serial number.

Many installations refer to their CPCs by a name that has meaning for them. If you supply an installation name in the CPC definition, SLM command responses and messages that have CPC information will contain the installation name for the indicated CPC.

From this panel...	select....
ThruPut Manager Main Lobby	4. Automation Services
Primary Options	2. Automation File Services
Automation File Services	1. Use an Existing File
Automation File pop-up	Enter to use the ACTIVE AF
Selection Menu	2. SLM - Service Level Manager
Policy Management	E to edit relevant policy
Policy Configuration Options	4. Automated Capacity Management
Automated Capacity Management (ACM)	3. Configure Central Processor Complex (LPAR Sets)

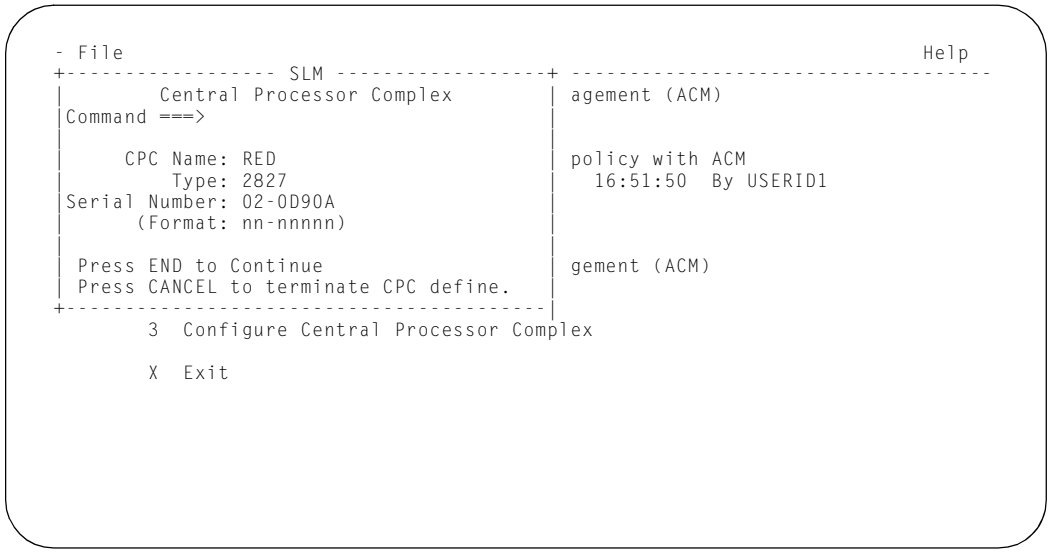

b. Define LPARs Running z/OS

Identify all of the LPARs on each CPC that are running z/OS so that SLM knows what LPARs are valid for inclusion in an LPAR Set and should be monitored.

After defining a CPC, a screen similar to the following is displayed.

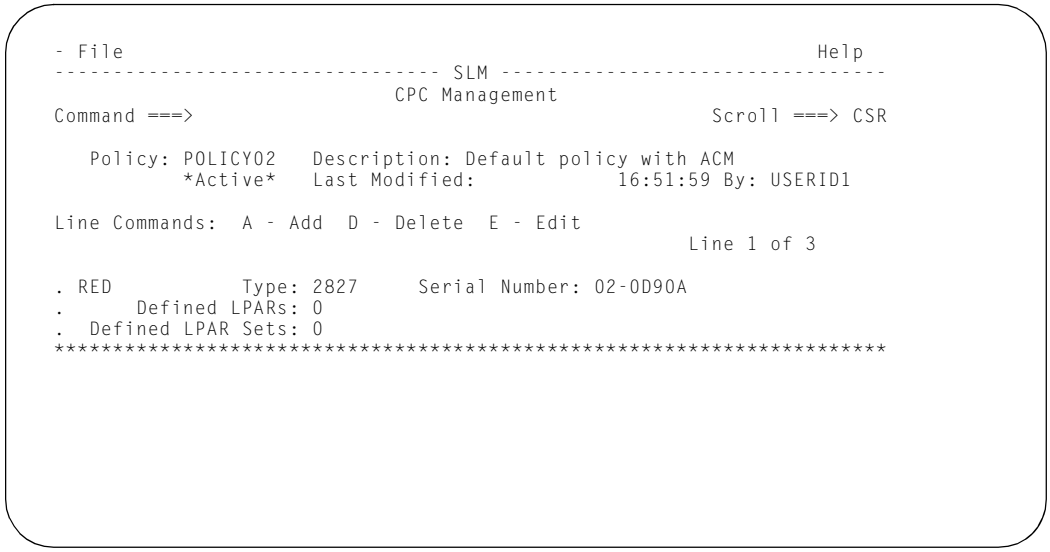

Enter the A-Add line command beside Defined LPARs to add another LPAR. A pop-up allows you to define the LPAR by entering the LPAR name.

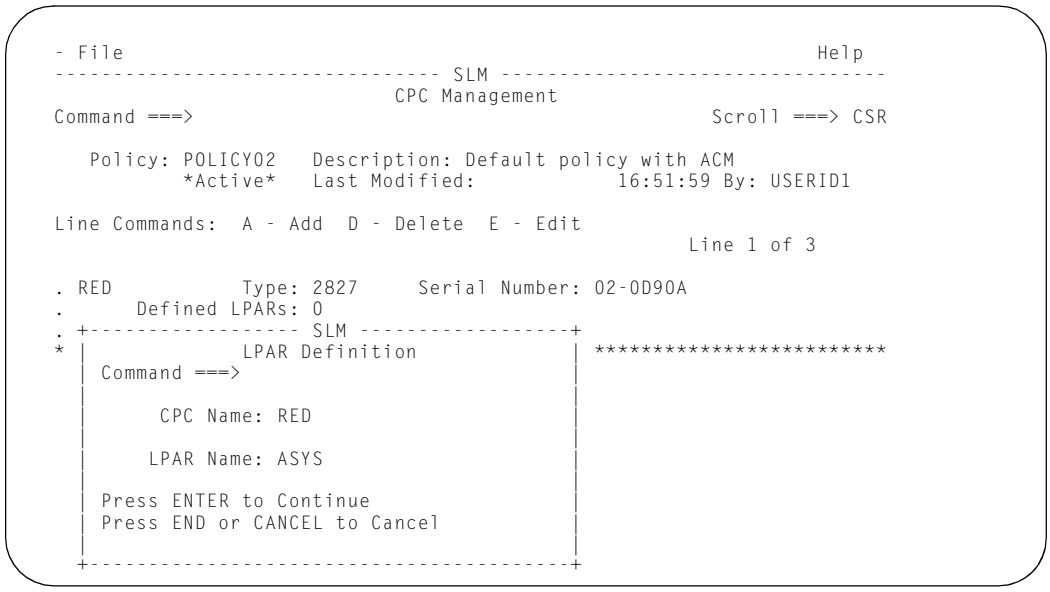

Continue to define all of the z/OS LPARs in each CPC.

c. Define LPAR Sets

Decide which LPARs should be grouped together with an associated limit for the Rolling 4-Hour Average (R4HA) in MSU/hr. For example all of the LPARs running IMS and DB2 might make up one group. An LPAR can be in more than one LPAR Set. Then, once you have defined your CPCs and LPARs, you can define the LPAR Sets on each CPC by choosing from the LPARs that you have defined, giving each LPAR Set a name and a limit. You may also designate it as a CMP LPAR Set and assign a CMP limit.

Enter the A-Add line command beside Defined LPAR Sets to add an LPAR Set. A pop-up allows you to define a new LPAR Set by supplying the name, R4HA limit and select the LPARs that belong to the set..

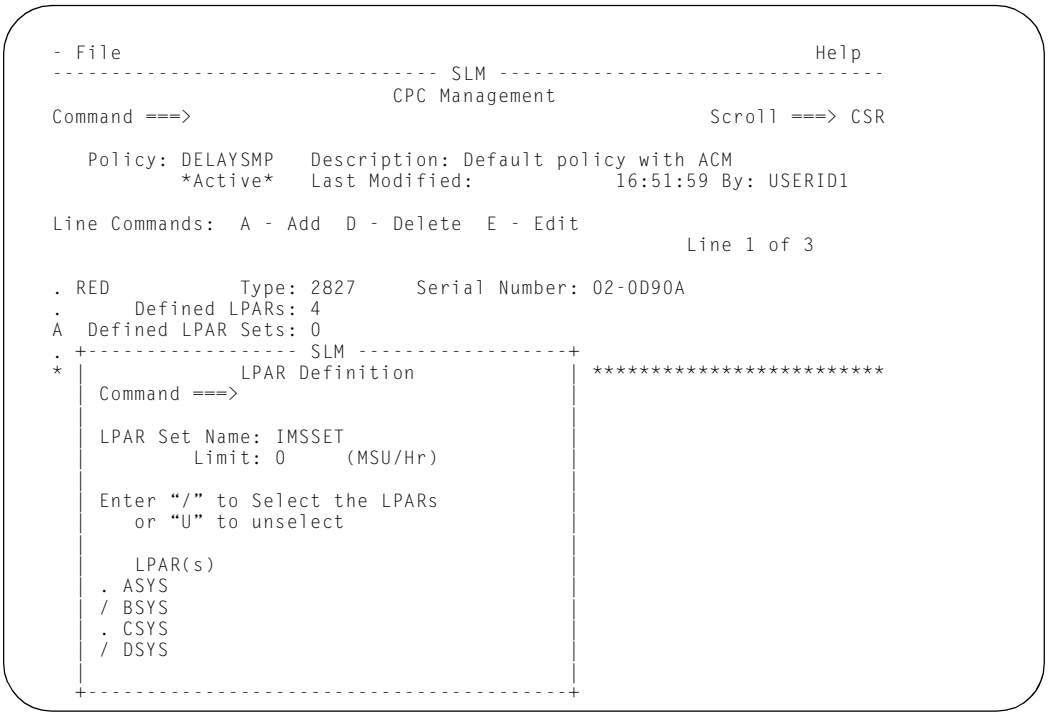

Continue to define all of the LPAR Sets that you have planned for each CPC.

If you make an error, you can edit the LPAR Set definitions using the E-Edit line and make the necessary changes.

> Note: if you are not sure what limit to set for a particular LPAR Set, set the limit to zero. After activating the policy, SLM will begin to monitor the usage and you can use the /SLM DISPLAY CPC operator command to display the current R4HA for the LPAR Set. Using that information, you update the limit for the LPAR Set to an appropriate value. This also applies to CMP LPAR Set limits.

3. Run ACM with Third Party Capping Software

If your installation is running third party capping software that is changing either the LPAR Defined Capacity or Capacity Group Limits, you will need to inform ACM using ACM Advanced Options..

From this panel...	select....
ThruPut Manager Main Lobby	4. Automation Services
Primary Options	2. Automation File Services
Automation File Services	1. Use an Existing File
Automation File pop-up	Enter to use the ACTIVE AF
Selection Menu	2. SLM - Service Level Manager
Policy Management	E to edit relevant policy
Policy Configuration Options	4. Automated Capacity Management
Automated Capacity Management (ACM)	2 Configure ACM Advanced Options

```
- File                                                          Help
----------------------------- SLM ------------------------------------
                        ACM Advanced Options
Command ===>

Policy: DELAYSMP     Description: Default policy with ACM
       *Active*      Last Modified:           16:51:59  By: USERID1

 MVS Solutions Support may request that you set one or more of the
following options if you are running third party software that changes
the LPAR Defined Capacity and/or Group Capacity Limits.

Please do not change these options unless requested.

     Enter "/" to enable option

         Ignore LPAR Defined Capacity
         Ignore LPAR Group Capacity
         Ignore capping status of LPAR

Press END to Continue
Press CANCEL to exit without saving any change(s).
```

In this case ACM should be instructed to ignore the limits being changed as it interferes with ACM's ability to determine how close the R4HA for the LPAR or Group is to the respective limit. In this case, LPAR Sets should be used to establish limits for use by ACM.

Please contact Compuware Customer Solutions assistance in setting the ACM Advanced Options.

4. Map ACM Service Classes to Capacity Levels

Map the Service Classes, reserved for ACM-specific use, on the following panel:

From this panel...	select....
ThruPut Manager Main Lobby	4. Automation Services
Primary Options	2. Automation File Services
Automation File Services	1. Use an Existing File
Automation File pop-up	Enter to use the ACTIVE AF
Selection Menu	2. SLM - Service Level Manager
Policy Management	E to edit relevant policy
Policy Configuration Options	1. Configure JESPLEX Characteristics
Automated Capacity Management (ACM)	1. WLM Service Class Definitions

```
- File                                                          Help
---------------------- SLM Definition Services ------------------------
                  WLM Service Class Definitions
Command ===>

   Policy: POLICY02  Description: Default policy with ACM
                     Last Modified:  12:34:54  By: USERID1

General Services                  Production Control Services
                                                  Premium: TMPS0

  Batch Importance 1: TMGS1          Production Importance 1: TMPS1
                   2: TMGS2          Production Importance 2: TMPS2
                   3: TMGS3          Production Importance 3: TMPS3
                   4: TMGS4          Production Importance 4: TMPS4
                   5: TMGS5          Production Importance 5: TMPS5
Automated Capacity Management (ACM)
 General Services
   Level 1: TMGSCL1   2: TMGSCL2   3: TMGSCL3   4: TMGSCL4   5: TMGSCL5

 Production Control Services
  Level 1: TMPSCL1    2: TMPSCL2   3: TMPSCL3   4: TMPSCL4   5: TMPSCL5

Press PF1 for information and guidelines on specifying Service Classes.
```

If you do not provide a Service Class for a capacity level, then the reduced service class constraint has no effect at that Capacity Level.

5. Specify Constraints by Production Importance

In our example Production Service jobs need to use reduced Service Classes at the following capacity levels:

- jobs with PI=5 will be constrained at capacity level 4
- jobs with PI=4 will be constrained at capacity level 3
- jobs with PI=3 will be constrained at capacity level 2

Jobs with PI=2 or 1 are not constrained at any capacity level..:

From this panel...	key....
ThruPut Manager Main Lobby	4 Automation Services
Primary Options	2 Automation File Services,
Automation File Services	1 Use an Existing File
AF pop-up	“Enter” to accept “active” file
Selection Menu	2 SLM - Service Level Manager
Policy Management	E to edit relevant policy
Policy Configuration Options	2 Configure Production Services
Configure Production Services	3 Automated Capacity Management
ACM Constraints	2 Constraints by Production Importance

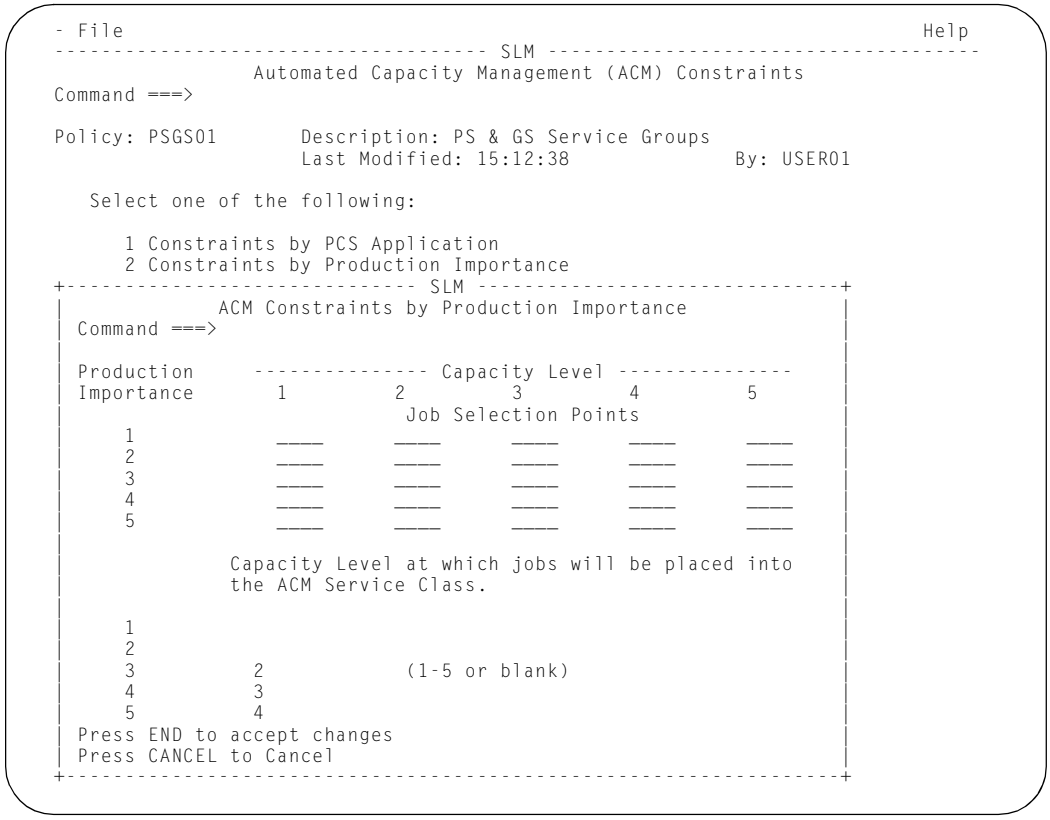

The Job Selection Point constraint is specified similarly to the one for the general constraint. (See the Usage Guide for details).

6. Specify Constraints for a Production Service Group

In our example, jobs assigned to the BNKEXPR PS Group need to be limited to:

- 16 jobs at capacity level 5
- 8 jobs at capacity level 4
- 6 jobs at capacity level 3
- 4 jobs at capacity level 2
- 1 job at capacity level 1

You can enter a slash opposite each capacity level for which you have a value:.:

From this panel...	key....
ThruPut Manager Main Lobby	4 Automation Services
Primary Options	2 Automation File Services,
Automation File Services	1 Use an Existing File
AF pop-up	“Enter” to accept “active” file
Selection Menu	2 SLM - Service Level Manager
Policy Management	E to edit relevant policy
Policy Configuration Options	2 Configure Production Services
Configure Production Services	2 Production Service Groups
Production Service Group Management	E-Edit for the relevant group
Production Service Group Definition	/ for Automated Capacity Management Constraints

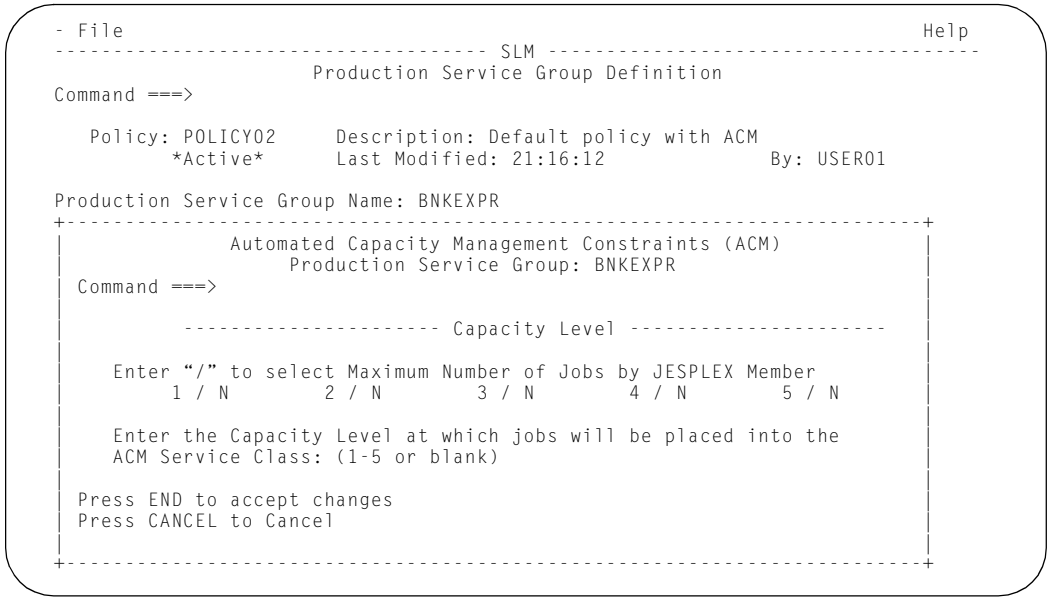

You step through each capacity level entering the appropriate value. Here's the first and last panel, with their values:

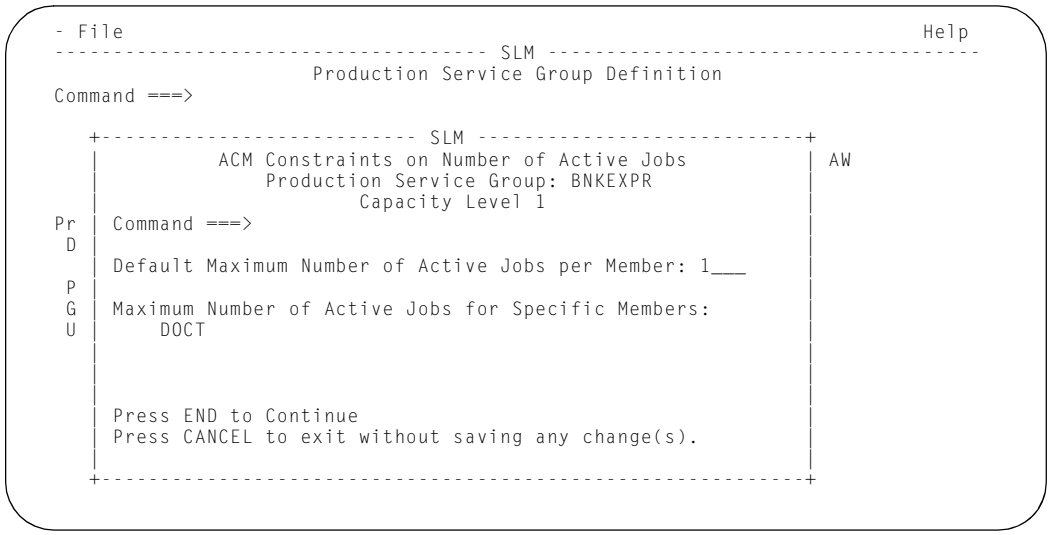

and

```
- File                                                                  Help
------------------------------------ SLM -------------------------------------
                     Production Service Group Definition
Command ===>

   +--------------------------- SLM ----------------------------+
   |          ACM Constraints on Number of Active Jobs          | AW
   |              Production Service Group: BNKEXPR             |
   |                    Capacity Level 5                        |
Pr | Command ===>                                               |
 D |                                                            |
   | Default Maximum Number of Active Jobs per Member: 16__     |
 P |                                                            |
 G | Maximum Number of Active Jobs for Specific Members:        |
 U |     DOCT                                                   |
   |                                                            |
   |                                                            |
   |                                                            |
   | Press END to Continue                                      |
   | Press CANCEL to exit without saving any change(s).         |
   |                                                            |
   +------------------------------------------------------------+
```

The Y on the following panel indicates a value has been specified..

```
- File                                                                  Help
------------------------------------ SLM -------------------------------------
                     Production Service Group Definition
Command ===>

   Policy: POLICY02     Description: Default policy with ACM
          *Active*      Last Modified: 21:16:12              By: USER01

Production Service Group Name: BNKEXPR
+-------------------------------------------------------------------------+
|              Automated Capacity Management Constraints (ACM)            |
|                   Production Service Group: BNKEXPR                     |
| Command ===>                                                            |
|                                                                         |
|           ---------------------- Capacity Level ----------------------  |
|                                                                         |
|    Enter "/" to select Maximum Number of Jobs by JESPLEX Member         |
|         1 / Y        2 / Y        3 / Y        4 / Y        5 / Y       |
|                                                                         |
|    Enter the Capacity Level at which jobs will be placed into the       |
|    ACM Service Class: (1-5 or blank)                                    |
|                                                                         |
| Press END to accept changes                                             |
| Press CANCEL to Cancel                                                  |
```

If they aren't already, enable these constraints on the ProductionServiceGroup Definition panel:.

```
- File                                                                  Help
------------------------------------ SLM ------------------------------------
                      Production Service Group Definition
Command ===>

   Policy: POLICY02     Description: Default policy with ACM
          *Active*      Last Modified: 21:16:12              By: USER01

Production Service Group Name: BNKEXPR
 Description: Helpful (not crucial) jobs

 Production Importance:    5   (1-5)
 Generate Job Summary:     Y   (Y/N)
 UDF Installation Panelid:

Service Times in the following line are expressed as mmm:ss
 2:00   to Target, 20:00  to Acceptable
                                                          Enabled
 Enter "/" to Update Constraints                           (Y/N)
          Constrain by Number of Active Jobs                 N
          Automated Capacity Management Constraints (ACM)    Y
```

You enter the Reduced Service Class constraint in a similar manner to that described for Production Importance constraints in Step 5.

7. Specify Constraints for a General Service Group

You specify the ACM constraints by updating the Control Center or Service Group from the General Service Group Management display. You can apply the three types of ACM constraints from this display: the JSP constraint is similar to that for general constraints (see the Usage Guide for details); the reduced service class constraint is similar to that in Step 5; the maximum number of jobs constraint is similar to that in Step 6.

Column C in the Service Group Management display shows if any constraint is in effect, whether normal ThruPut Manager AE constraints or ACM ones. A Y in the Constraints column of a Control Center applies to all of the Service Groups (Types) within that Control Center.

```
- File  View                                                            Help
------------------------------------ SLM ------------------------------------
                       General Service Group Management
Command ===>

   Policy: POLICY02   Description: Default policy with ACM
          *Active*    Last Modified: 21:16:12  By:        USER01

Line Commands: C - Copy   D - Delete  E - Edit  M - Modes

--Service Group--
Control             B Aging  --------Service Times---------            Service
Center-- Type----   I Limit-Target-Accept-Critical C  Description------- Class
. BNK                                              Adhoc banking jobs
.        DBUPLOAD   5   A    15:00  45:00           Refresh Test DBs   TMGS5
.        DEV        3   A     5:00  10:00         Y Compiles, etc      TMGS3
.        QA         2   C     5:00  20:00 30:00     Pre-prod testing   TMGS2
********************************************************************************
```

8. Enable ACM Constraints

ACM is enabled at two levels: the constraints for each Service Group can be enabled and disabled independently; and, the overall ACM feature can be enabled or disabled. When the overall feature is enabled, the PI-based constraints and the enabled constraints from Production and General Service

Groups are applied as the capacity level changes. If the overall feature is disabled, no ACM constraints are applied. Use the following panel to enable the overall ACM feature..

From this panel...	key....
ThruPut Manager Main Lobby	4 Automation Services
Primary Options	2 Automation File Services,
Automation File Services	1 Use an Existing File
AF pop-up	“Enter” to accept “active” file
Selection Menu	2 SLM - Service Level Manager
Policy Management	E to edit relevant policy
Policy Configuration Options	4 Automation Capacity Management
Automation Capacity Management (ACM)	1 Configure Automation Capacity Management (ACM)

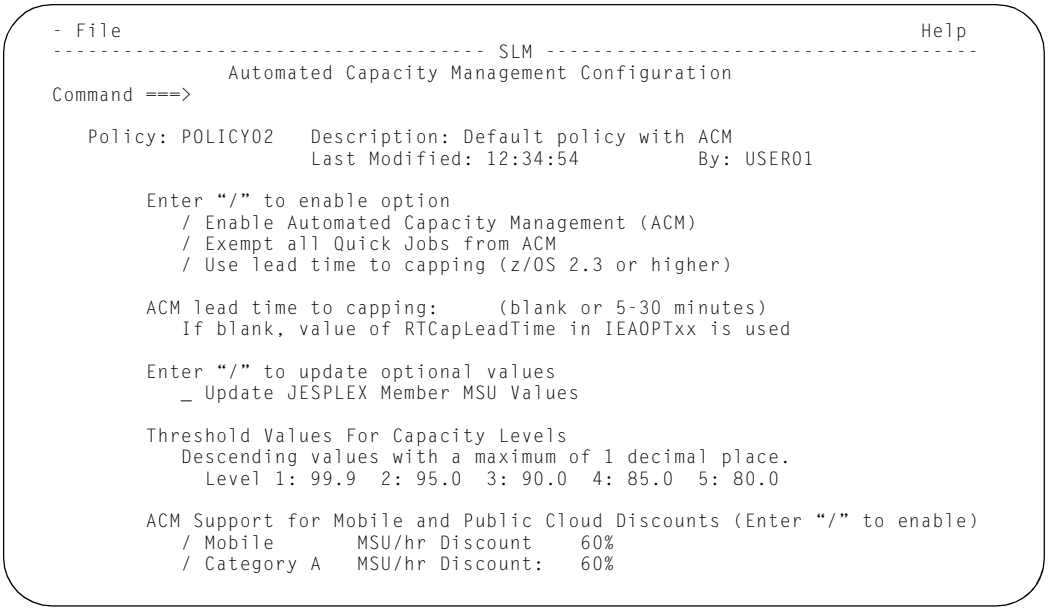

Unless you enable ACM, it does not take effect, and ThruPut Manager AE processes jobs just as if this elective functionality were not available.

9. Decide if Quick Jobs are Constrained

Use the above screen to exempt quick jobs from being subject to ACM constraints.

10. Enable ACM Support of WLM Estimated Time to Capping

Select the "Use lead time to capping" option to enable ACM support for the estimated time to capping of the LPAR or LPAR Group, made available by WLM in z/OS 2.3 and higher releases.

If selected, ACM will use the lead time value supplied in the RTCapLeadTime parameter in the active IEAOPTxx member of the z/OS PARMLIB to compare against the time estimate. If the estimate is less that the lead time, ACM will set the current Capacity Level to 1, which will activate the maximum constraints on batch workload specified by the installation in the active SLM Policy.

If desired, you can also provide a value on the panel ("ACM Lead time to capping"), which ACM will use instead of the value of RTCapLeadTime. Note that the IBM-supplied default value for RTCapLeadTime is 0.

11. Enable Discounting of Mobile and/or Public Cloud Workloads

Select “Mobile” and/or “Category A” to enable ACM automatic discounting of workloads classified by WLM as MOBILE and/or CATEGORYA (public cloud) from the calculated R4HA (rolling 4-hour average) and 5 minute interval consumption rate for every CMP LPAR Set.

At present, both of IBM’s Mobile Workload Pricing (MWP) and Z Workload Pricing for Cloud (zWPC) pricing options specify a discount percentage of 60% which is therefore the default value used in both cases by ACM. If the discount percentages change in the future, they may be modified in the “MSU/hr Discount” fields on the panel.

12. Specify Capacity Limits (optional)

The “Update JESPLEX Member MSU Values” option enables you to specify capacity limits for individual LPARs, in terms of Millions of Service Units (MSUs per hour). This function applies only if your datacenter has not specified LPAR Defined Capacity and/or LPAR Group Limit using the Hardware Management Console (HMC) or if you are using a third party capping product and have told ACM to ignore the Defined Capacity and Group Limit.

When you choose to Update JESPLEX Member MSU Values, a pop-up appears:.

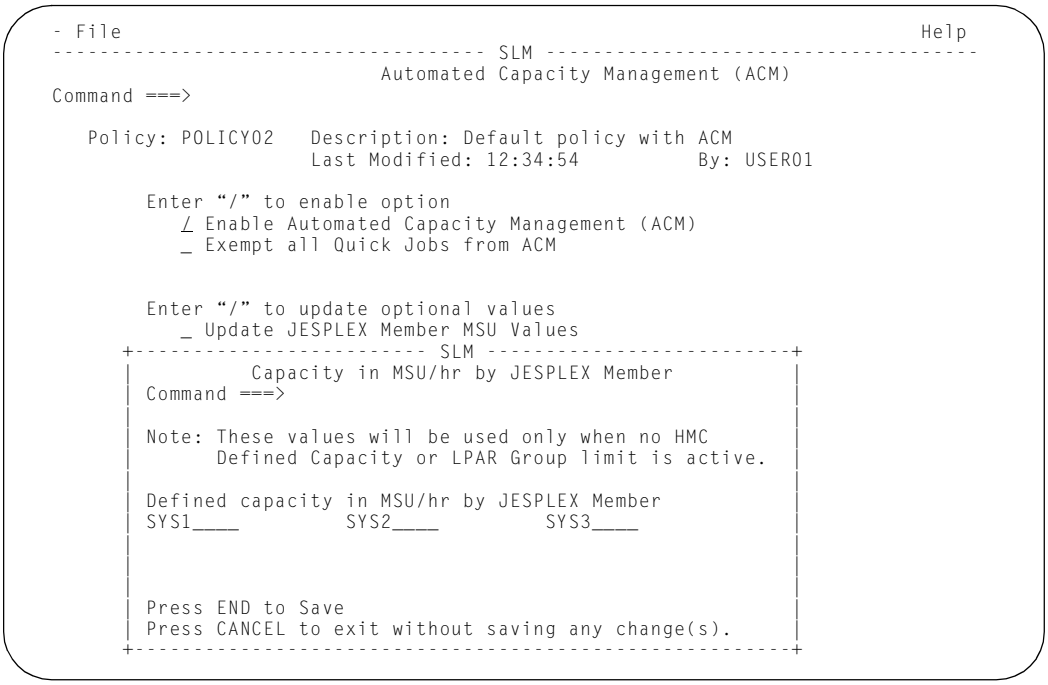

13. Adjust Percentages for Capacity Levels 2 to 5 (optional)

The “Threshold Values For Capacity Levels” option allows you to change all but the highest default percentages at which the various Capacity Levels take effect. In calculating these capacity thresholds, ACM takes into account the short term usage trend and whether or not the LPAR is currently soft capped.

You must keep Capacity Level 1 at 99.9%, but can change the values for Capacity Levels 2 through 5. This would be appropriate if, for example, your datacenter has very high capacities, so that a single

percentage point represents a considerable amount of resources. In this case you might want to compress the total range of values, starting with Level 5 at, say, 90%.

14. Add JAL to Exempt Specific Jobs

You can exempt a job from Automated Capacity Management control and processing, even if its Service Group has ACM constraints specified for it, by specifying in JAL:

```
SLM SET NO_CAPACITY_MANAGEMENT
```

15. Activate Your Updated Policy

Once you have specified your Automated Capacity Management constraints, you can activate the new Policy. Remember to refresh your JAL if modified.

Note that, if you run Generate Report for the Policy, the ACM Service Classes are listed as well as the ACM constraints.

Chapter 3
Run with ACM Constraints

ThruPut Manager AE, with ACM enabled, automatically takes actions as Capacity Levels change.

Automatic Runtime Behavior

Changes in the Capacity Level are recorded in the System Log with the following messages:

```
DTM8161I memname AT CAPACITY LEVEL n: USAGE IS increasing/decreasing

DTM8162I memname HAS REACHED FULL CAPACITY

DTM8163I memname CAPACITY CONSTRAINTS RELIEVED
```

At selection time all enabled constraints, whether from ThruPut Manager AE or Automatic Capacity Management, must be satisfied before a job is selected.

The WLM Service Class assigned to a job depends on the current Capacity Level and the constraints you have specified. If the R4HA is below Capacity Level 5, or the Capacity Level is below the one you have specified to put these jobs into an ACM Service Class, the Service Class assigned is the one normally used, i.e., the one based on the job's importance.

A job with a Reduced Service Class constraint is placed in the ACM low-resource Service Class associated with the current Capacity Level. During job execution, ACM changes the Service Class of a job, as required, to that associated with the current Capacity Level.

A job with a Maximum Number of Jobs constraint is selected normally until the Capacity Level of the constraint is reached and the maximum number of jobs belonging to this Service Group (or Control Center) is already being processed. Once the number of jobs falls below the maximum or the Capacity Level changes to reflect a more lenient constraint, the job is again eligible to be selected.

A job with a Job Selection Point constraint is selected normally until the Capacity Level of the constraint is reached and the job's position in the queue remains below the job selection point specified for it. Once the job reaches the job selection point or the Capacity Level changes to reflect a more lenient constraint, the job again becomes eligible to be selected.

Changes to the SLM Dialog

In the SLM dialog, you see ACM information in the pop-up window accessed from the JOB DISPLAY panels. The Estimated Time to Selection field may have Indeterminate or Indeterminate-Delayed for ACM-constrained jobs. Using the I-Information line command shows greater detail, for example:.

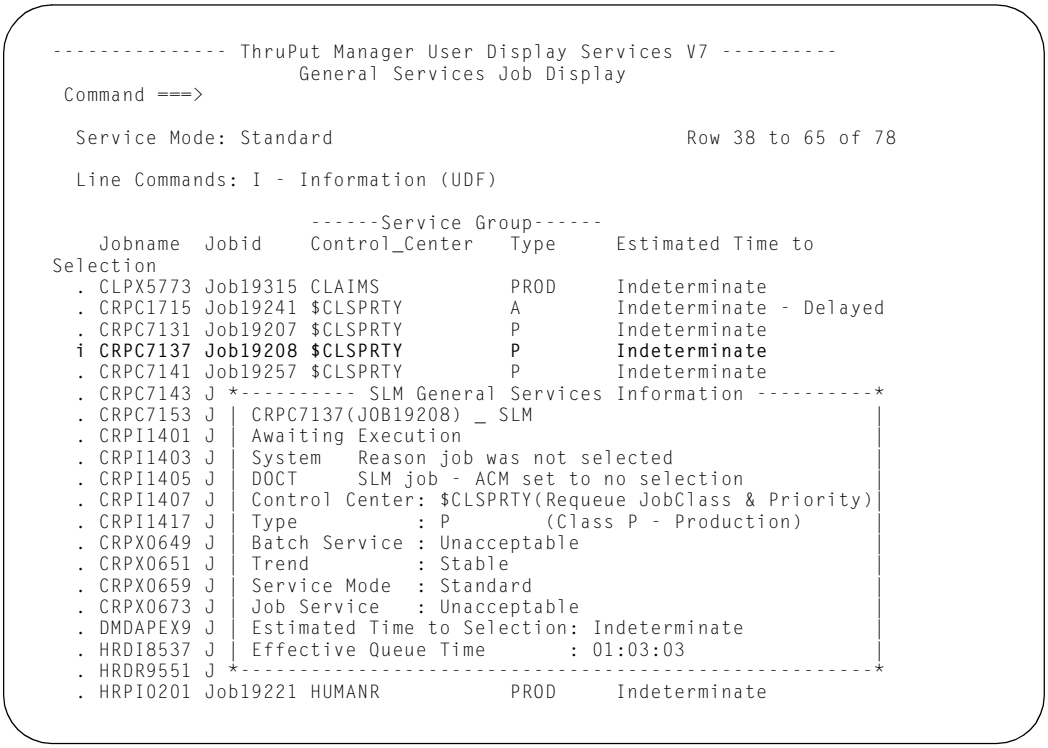

Changes to Other Dialogs

ThruPut Manager AE supplements SDSF, IOF, and EJES information. The SDSF Status

column may have new values, such as:

- SLM Cap - Sel Pt - this job cannot run because it is currently below the selection point required to be selected when the system is at the current Capacity Level.
- SLM Cap - No Sel - this job cannot run because it is currently not allowed to be selected when the system is at the current Capacity Level.
- SLM Cap - Max Job - this job cannot run because the Service Group assigned to this job is already running the maximum number of jobs specified for the current Capacity Level.
- SLM Cap - Max CC - this job cannot run because the Control Center assigned to this job is already running the maximum number of jobs specified for the current Capacity Level.

Note: Both the jobname and jobnumber must be visible on the panel in order for the extended TM data to appear in the Status column on both the Status and Input panels.

The SRVCLASS column shows when the job is executing in an ACM Service Class..

```
 Display  Filter  View  Print  Options  Help
-------------------------------------------------------------------------------
SDSF INPUT QUEUE DISPLAY ALL CLASSES                        LINE 27-60 (66)
COMMAND INPUT ===>                                           SCROLL ===> CSR
NP   JOBNAME   JobID     Owner    Prty C Status             SrvClass PrtDest
     CRPC7137 JOB21587   CA750NL     2 1 SLM Heavy Load     TMGS3    LOCAL
     CRPC7149 JOB21532   CA750NL     8 1 SLM Cap - Sel Pt   TMGS3    LOCAL
     CRPC7153 JOB21466   CA750NL     8 1 DCS delay          TMGS3    LOCAL
     CRPX0659 JOB21513   CA750NL     8 1                    TMSLMBI3 LOCAL
     CRPX1939 JOB21516   CA750NL     9 1 SLM Cap - Sel Pt   TMGS3    LOCAL
     CRQC719  JOB21185   CA750NL     9 1 DCS delay          TMGSCL2  LOCAL
     CRXX4724 JOB21588   CA750NL     6 Y                    BATCH    LOCAL
     CRXX4735 JOB21190   CA750NL     2 1                    TMGSCL2  LOCAL
     HRDI8531 JOB21202   CA750NL    10 1 SLM Cap - Max CC   TMGS5    LOCAL
     HRDR9549 JOB21205   CA750NL    10 1 SLM Cap - Max CC   TMGS5    LOCAL
     HRPI0201 JOB21197   CA750NL    12 1 SLM Cap - Max CC   TMGS3    LOCAL
     HRPI3521 JOB21165   CA750NL    12 1 SLM Cap - Max CC   TMGS3    LOCAL
     SO42CB51 JOB21002   CA750NL     2 1                    TMGSCL2  LOCAL
     SO42CB52 JOB21156   CA750NL     2 1                    TMGSCL2  LOCAL
     SPO72CCE JOB20411   CA750NL     9 D HELD ThruPut Mgr   BATCH    LOCAL
     SPO72CCH JOB21304   CA750NL     2 1                    TMGSCL2  LOCAL
     CRXX4233 JOB21581   CA750NL     6 Y                    BATCH    LOCAL
     WLDC903  JOB21450   CA750NL     7 1 DCS delay          TMGSCL2  LOCAL
     WLDC909  JOB21528   CA750NL     9 1 SLM Cap - Max CC   TMGS4    LOCAL
     WLDI872  JOB21518   CA750NL     2 1                    TMGSC2   LOCAL
     WLPR441A JOB21564   CA750NL     6 1 SLM Cap - Max CC   TMGS4    LOCAL
```

Using the TMUSER command for a job yields more information:.

```
--------------- ThruPut Manager User Display Services V7 ---------------
 Command ===>                                          SCROLL ===> CSR
 SDSF INPUT QUEUE DISPLAY ALL CLASSES                  LINE 27-60 (66)
 COMMAND INPUT ===>                                    SCROLL ===> CSR
 NP   JOBNAME  JobID    Owner    Prt *---------- Job List Display -------*
      CRPC7137 JOB21587 CA750NL      | _ SLM                             |
      CRPC7149 JOB21532 CA750NL      | _ SLM                             |
      CRPC7153 JOB21466 CA750NL      | _ SLM DC                       H  |
      HRDI8531 JOB21202 CA750NL    1 | _ SLM JB                          |
      HRDR9549 JOB21205 CA750NL    1 | _ SLM                             |
      HRPI0201 JOB21197 CA750NL    1 | _ SLM                             |
      HRPI3521 JOB21165 CA750NL    1 | _ SLM                             |
      *-------------- SLM General Services Information --------------*    |
      | WLPR441A(JOB21564) _ SLM                                    |    |
      | Awaiting Execution                                          |    |
      | System   Reason job was not selected                        |    |
      | DOCT     SLM job - ACM is at maximum jobs for Control Center | H  |
      | Control Center: WHOLELIF (Whole Life Insurance Division)    |    |
      | Type          : A        (Class A jobs)                     |    |
      | Batch Service : Unacceptable                                |    |
      | Trend         : Stable                                      |    |
      | Service Mode  : Standard                                    | H  |
      | Job Service   : Beyond Target                               |    |
      | Estimated Time to Selection: Indeterminate                  |    |
      | Effective Queue Time       : 00:10:45                       |    |
      *-------------------------------------------------------------*    |
     WLPI0209 JOB21517 CA750NL      | _ SLM                             |
 i   WLPR441A JOB21564 CA750NL      | _ SLM                             |
     WLPR441A JOB21586 CA750NL      | _ SLM                             |
     WLPR441B JOB21582 CA750NL      | _ SLM                             |
     WLPR441C JOB21583 CA750NL      | _ SLM                             |
```

The TMUSER synchronized window (Job List Display) has the same ACM information if you drill down on SLM. In addition, choosing SLM Job Summary and scrolling through its pop-up window shows that the Job Summary includes JOB selection was affected by capacity management.

Enhanced ThruPut Manager AE Commands

ThruPut Manager AE commands provide additional information when ACM is enabled.

```
/SLM DISPLAY
DTM8213I SLM SUMMARY
POLICY                    : POLICY02
SERVICE MODE              : STANDARD
SLM SERVER ON MEMBER      : DOCT
GS OPERATOR CONSTRAINTS   : NONE
BATCH SERVICE             : UNACCEPTABLE
CAPACITY MANAGEMENT       : CONFIGURED
```

Use /SLM Display JESPLEX to see the current R4HA, Capacity Level, CPU usage in the last five minute interval and the percentage of time during that interval when the system was capped.

If the LPAR is part of an LPAR Group, you also see the LPAR Group name, the group R4HA, the group usage in the last five minutes, and the LPAR's current group share. Note that each LPAR Group member does its own calculation of LPAR Group usage, so they may differ slightly from each other. This is normal.

If the LPAR is part of one or more LPAR Sets, you also see a line for each LPAR Set name, the R4HA for the set and the usage by the set in the last five minutes. If the LPAR Set is a CMP LPAR Set, you will see an additional line showing the aggregate R4HA and 5 minute interval usage.

If mobile or public cloud discounting has been activated in the SLM Policy and any workload of the appropriate classification is present, the R4HA and interval usage of the workload in each LPAR Set is also displayed.

```
/SLM D JESPLEX
DTM8219I SLM JESplex Status
Member: SUP1   State: Member Inactive
Member: SUP2   State: Member Inactive
Member: SUP3   State: SLM Server Active
  BC12      Type: 2828  Serial: 02-C0A32
  Logical Processors  Active: 2    Parked: 0
  Capacity Management at 2015.292 14:30:29
    SLM Capacity Level: 1
    SUPP3    * 4h Avg: 3.01      Intvl: 0.50       Cap: 100%
      Avg jobs limited by Service Class: 0.0
    SUPPSET    4h Avg: 3.35      Intvl: 0.76
  Constraints:
    General Services: Color(Blue)
Manager: Enter STOP or new time interval
```

Use /SLM Display CPC to display CPC type and serial number as well as the current R4HA for LPARs and LPAR Sets which have been defined in the active SLM Policy.

```
/SLM DISPLAY CPC
DTM8508I SLM CPC Status
CPCA      Type: 2828  Serial: 02-C0A32
LPARs:
  SUPP1     4h Avg: 0.00      Intvl: N/A
  SUPP2     4h Avg: 0.29      Intvl: 0.23
  SUPP3     4h Avg: 3.05      Intvl: 0.89
  ETP1      4h Avg: 0.82      Intvl: 0.76
  ETP2      4h Avg: 0.76      Intvl: 0.74
LPAR Sets
  SUPPSET   4h Avg: 3.34      Intvl: 1.12  Lmt: 4
  ETPSET    4h Avg: 1.58      Intvl: 1.50  Lmt: 5
```

The /SLM Display SC command shows the SLM Service Classes currently in use, including the ACM Service Classes, if any.

```
/SLM DISPLAY SC ALL
DTM8233I SLM Service Class Display
...
TMGSCL1   Per:1   PI:1.8   Dly:54%   ACPU:0.00%  Clr:Y  AS:17.0
TMGSCL2   Per:1   PI:N/A   Dly:N/A   ACPU:0.00%  Clr:B  AS:0.0
TMGSCL3   Per:1   PI:N/A   Dly:N/A   ACPU:0.00%  Clr:B  AS:0.0
TMGSCL4   Per:1   PI:N/A   Dly:N/A   ACPU:0.00%  Clr:B  AS:0.0
TMGSCL5   Per:1   PI:N/A   Dly:N/A   ACPU:0.00%  Clr:B  AS:0.0
TMGS1     Per:1   PI:0.6   Dly:4%    ACPU:1.17%  Clr:B  AS:2.0
TMGS2     Per:1   PI:N/A   Dly:N/A   ACPU:1.17%  Clr:B  AS:0.0
TMGS3     Per:1   PI:N/A   Dly:N/A   ACPU:1.17%  Clr:B  AS:0.0
TMGS4     Per:1   PI:N/A   Dly:N/A   ACPU:1.17%  Clr:B  AS:0.0
TMGS5     Per:1   PI:N/A   Dly:N/A   ACPU:1.17%  Clr:B  AS:1.0
```

The /SLM Display SG command includes ACMJobs when the ACM constraints for maximum number of jobs is in effect for that General Service Group.

```
/SLM DISPLAY SG BNK DEV
DTM8220I SLM Service Group Status
Control Center: BNK
  JESplex Constraints: None
  Constraints by JESplex Member:
    DOCT  ACMJobs
Type: DEV
  JESplex Constraints: AgeLmt
  Constraints by JESplex Member:
    DOCT  Color(Blue)
```

Full syntax of these commands can be found in the *Command Reference Guide*.

Running With ACM - The "New Normal"

Running ThruPut Manager AE with ACM enabled is like the "New Normal" described in the *ThruPut Manager AE Concepts Guide*, with a few additions due to ACM.

Soft capping, (if you've enabled it), occurs much less frequently and for shorter durations. The most important batch work continues to run, and the online doesn't seem to be affected.

Savings start to appear in budget reviews due to lowered capacity limits, and reduced software MLC charges.

Users notice their important jobs continue to meet their targets, though other, less important jobs sometimes appear to have longer system residencies. SDSF information shows ACM information in the Status column and pop-up windows. A messages appears in the Job Summary Report:

```
Job Selection was affected by capacity management
```

Technical staff can readily monitor the R4HA with a simple ThruPut Manager AE command, see how it compares to your capacity limits, and get recent trend information. There are new messages that ThruPut Manager AE issues when the Capacity Level changes on an LPAR.