




The Mainframe Software Partner For The Next 50 Years

ThruPut Manager PCS Concepts Guide

Release 18.02



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

ThruPut Manager Customer Support

<https://go.compuware.com/>

This document and the product referenced in it are subject to the following legends:

Copyright 2019 Compuware Corporation. All rights reserved. Unpublished rights reserved under the Copyright Laws of the United States.

U.S. GOVERNMENT RIGHTS-Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in Compuware Corporation license agreement and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (OCT 1988), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable. Compuware Corporation.

This product contains confidential information and trade secrets of Compuware Corporation. Use, disclosure, or reproduction is prohibited without the prior express written permission of Compuware Corporation. Access is limited to authorized users. Use of this product is subject to the terms and conditions of the user's License Agreement with Compuware Corporation.

Adobe® Reader® is a trademark of Adobe Systems Incorporated in the United States and/or other countries.

All other company or product names are trademarks of their respective owners.

Introduction

Summary of Changes

V1802-7118 (April 2019)	<ul style="list-style-type: none"> • No changes
V1802-7117 (January 2019)	<ul style="list-style-type: none"> • No changes
V1802-7116 (October 2018)	<ul style="list-style-type: none"> • No changes
V1802-7115 (July 2018)	<ul style="list-style-type: none"> • No changes.
V1802-7114 (April 2018)	<ul style="list-style-type: none"> • This is the base manual for ThruPut Manager 18.02.

About This Manual

Production Control Services (PCS) is a feature of ThruPut Manager Automation Edition. This manual describes the concepts and ideas that form its approach to managing your CA 7 scheduled workload. It include managing the critical path, predicting when jobs will be late and mitigating that situation based on the importance of the work. This is the starting point for anyone with an interest in PCS.

This *PCS Concepts Guide* includes a chapter, [PCS Core Features](#), which includes information applicable to any PCS implementation. The second chapter, [PCS Advanced Highlights](#), answers the question, "What else can I do with PCS?" The advanced features are mostly independent of one another and are built on the core features. For each one, an overview and where to find additional information is provided.

The *PCS Usage Guide* leads you through creation of the constructs that every PCS implementation must build. It also shows you how operations and your end-users interact with PCS at runtime.

Reference manuals are available and applicable while implementing PCS functionality. They supply the details on rules, commands, messages, and administrative or environmental tasks.

This guide assumes the reader is familiar with the *TM AE Concepts Guide* and the *TM AE Usage Guide*.

Contents

Introduction	3
Summary of Changes	3
About This Manual	3
Chapter 1 PCS Core Features	7
ThruPut Manager AE and PCS	7
Definition of "PCS Core Features"	7
PCS Description of Your Workload	7
Critical and Urgent Paths	8
"Lateness"	9
Milestone Events	9
Production Importance	9
Introduction to PCS Constructs	9
Application	9
Production Cycle	10
Experience File	10
Battle Plan	10
Installation List	11
Alerts	11
Production Service Groups Considerations	11
PCS Workload Management Process	12
Runtime View	13
Operator Control and Intervention	14
Implementation Considerations	15
Typical PCS Day	16
Conclusion	16
Chapter 2 PCS Advanced Highlights	17
Prerequisites	17
Advanced Highlights	17
Automated Capacity Management	17
Quick Jobs	17
Demand Job List	18
External Process List	18
Appendix A Production Service Groups	19
Overview	19
Comparison with other jobs in the PS queue	19

Chapter 1

PCS Core Features

This chapter explains the core features and facilities of PCS.

ThruPut Manager AE and PCS

ThruPut Manager AE (Automation Edition) is a software product that optimizes the total z/OS JES2 batch workload, by automating and streamlining batch job processing, providing service-oriented automation and tools, and reducing manual intervention previously required.

The influence of ThruPut Manager AE extends from submission of a job to the end of execution of that job. It makes sure the jobs are processed efficiently and automatically, reacting to processing shortages and other problems as they arise. Further, ThruPut Manager AE has influence over the entire batch workload. Accordingly, it is able to minimize any interference of your critical production applications from non-production or lower priority production jobs.

PCS extends this service model to a CA 7 scheduled workload by taking into account the relationship amongst the scheduled jobs, the desired completion time and the critical path. It understands this relationship and deploys the rest of the ThruPut Manager features to optimize the scheduled workload.

Definition of "PCS Core Features"

We have identified the subset of PCS features that is needed to extend your CA 7 goals into the ThruPut Manager scope, and that is applicable to any datacenter using PCS. This subset is called the *PCS core features*, and its setup and deployment is a *PCS core implementation*.

There are several other, mostly independent, features that are available to address issues particular to your datacenter. These are referred to as *PCS advanced features*, and assume you have completed the PCS core implementation. These advanced features are introduced in [PCS Advanced Highlights](#).

PCS Description of Your Workload


The mission of ThruPut Manager is to automate and optimize the running of your z/OS JES2 batch workload. PCS is the part of ThruPut Manager that understands that some workloads are scheduled (by CA 7) and can make decisions about individual jobs as part of a scheduled workload. PCS understands where the critical path lies and which jobs are part of it, when the critical path shifts to another path, when a job is expected to complete and when a job is projected to be late.

PCS does more than "know" and "understand". PCS is proactive. For instance, deploying the ThruPut Manager engine, it favors higher importance jobs over lower importance jobs, and, for late jobs, changes Service Classes to one with higher WLM service goals. PCS provides comprehensive automation to manage your scheduled batch production workload, in an optimal manner.

Because of its unique perspective, PCS can present an integrated display that includes CA 7, JES2 and ThruPut Manager information, giving the operations staff the information they need to know to ensure ThruPut Manager is handling the scheduled workload in the optimal manner.

PCS keeps a detailed view of your scheduled workload. It interrogates the CA 7 database and builds up a picture of each application, the job trees within each application, the job paths within each job tree, and the jobs within each job path. Over time it builds a history or "experience" data for each job

that contributes to the expectations for how long a path, tree and application is expected to run. It is very good at predicting when an application "will be late" early in the game.



Here's an application with 2 job trees, one of which has a subtree. The job trees are different shades of blue; the subtree within the darker blue tree also has a drop shadow. The solid and dashed red lines each trace a different path. There are 6 paths and 16 jobs in this application.

PCS has four sources of information:

- the CA 7 database: PCS captures the application-tree-path-job dependency structure; due out times, if entered; and, the actual jobs running on a specific day.
- the Application Management Database (AMD): Populated by the installation, you can add values to indicate the importance of the work, when it should be considered late, whether it should have early analysis, when to raise an alert, and another setting to control how PCS views and processes your workload.
- the SLM Policy: Populated by the installation, you provide information pertaining to the Production Services (PS) queue and Services Classes to be used for this workload. The PS queue and service classes are also used by the Production Service Groups, if deployed in your installation.
- the experience data: Captured automatically, PCS aggregates and uses this information to predict elapsed times, completion times, and if a job will be or is late.

Of course the existing sources that ThruPut Manager uses are still relevant when the job is processed by ThruPut Manager: job profile determined in the analysis process, JAL statements, JECL statements, binding and limiting agents and so on.

Critical and Urgent Paths

PCS identifies both the critical path(s) for the workload and urgent paths for a tree or for a job. Both types of paths gives you information about your workload.

A critical path in project management and general network theory is the path with the longest duration (the summation of its dependent parts) needed to complete the project or transverse the network. It gives a theoretical minimum completion time which cannot be reduced further.

In PCS, the critical path is the path, within all the applications, with the least rerun factor (a ratio of how long a job would take to rerun to how much time is available to rerun it) when taking into account both predecessor and triggering relationships. More than one critical path is identified if their rerun factors are very close to one another.

The urgent path is identified for each tree: it's the path with the least rerun factor within the tree, taking into account its triggering relationships. The urgent path is also identified for each job: it's the triggering path with the least rerun factor that the job is part of. A urgent path for a job need is not necessarily an urgent path of a tree.

"Lateness"

PCS uses *acceptable time* to indicate the date and time by which a job needs to complete to be considered *not late*. PCS uses a proprietary algorithm to predict a job will be late, involving expected run time (from the experience data), slack time (the time between jobs in a path), and the progress of upstream jobs. From this it deduces if a path, tree or application would therefore be late.

When looking at a job in jeopardy of being late, PCS can tell you the *culprit job* (an upstream job with the initial problem) and any other affected jobs.

Milestone Events

Milestone jobs can be created by specifying the acceptable time for a job, path or tree. PCS adjusts the acceptable time, slack time and rerun factor of upstream jobs to make sure they are handled with this 'milestone' in mind. Alerts can be triggered by any milestone event.

Production Importance

PCS uses an aging algorithm based on the urgency of the work to determine the next job to select. (It never artificially delays work that would otherwise be "early".)

When the system can't deliver all the jobs on time, it takes into consideration the importance of the workload and delivers the most important earlier. Importance is indicated by setting values from 1 to 5 of production importance in the PCS dialog, where 1 is the most important.

Introduction to PCS Constructs

PCS users deal with the following ideas and constructs to specify and monitor the PCS-pertinent information.

Application

Applications correspond to the CA 7 System name and the list of applications is determined by interrogating the CA 7 database. PCS also determines the tree-path-job structure for each application.

Though PCS has initial values for the defaults and options that apply to applications, you can tailor these to your needs with overrides for specific applications, trees, paths, or jobs.

You can add *alerts* based on triggering events at the application, tree, path, or job level.

Applications are in either managed or monitored mode. If an application is in managed mode PCS manages the jobs in the application by directing SLM how to optimize the jobs. If an application is in monitored mode PCS displays information and processes alerts for the relevant jobs, but does not

interfere with the actual processing of the job (it continues to be processed as it would without PCS). An application can also be neither managed nor monitored, in which case PCS only collects experience data for it.

PCS stores this information (and other information) in the *Application Management Database (AMD)*. Generally you have one AMD per CA 7 instance.

Production Cycle

A *production cycle* is a time period during which PCS considers CA 7 submitted jobs for PCS processing. This set of jobs is dependent on the output of the PCS Forecast Collection Utility job. This utility accesses CA 7 schedule data for the next few days and determines which jobs are scheduled for the upcoming production cycles.

A common and simple production cycle setup is one cycle daily, say 7:00 a.m. to 6:59 a.m. the next day. But other options are possible. For example, you might have multiple production cycles in a day that align with your shift schedules.

How does one choose a production cycle's start and end times? Consider when scheduled batch runs. When is the natural time to pick up CA 7 changes? When is a natural time for the default late time to be? One installation with frequent CA 7 changes start their production cycles in the late afternoon, after the PCS Forecast Collection Utility has been run (~6:00 p.m.). They chose an end time of early morning (~7:00 am) to establish an appropriate late time for the overnight work. They have a second production cycle which runs from 7:00 am to 6:00 p.m. to cover daytime scheduled work. Other installations consider their shift times, the time zones they cover and other patterns of work.

Production cycles are designed to repeat on a weekly basis. However, one-time production cycles are possible for exceptional circumstances when you have a change that needs to override the usual setup.

PCS provides features to manage jobs with the same name that repeat within the same production cycle, each with different acceptable times.

Though production cycles may endure for up to 48 hours, we assume they are 24 hours or less in the "core" discussion. We also assume that there are no gaps between production cycles. Other variations are discussed in Crossing Production Cycle Boundaries chapter in the *PCS Usage Guide*.

Experience File

PCS collects information on each job as it completes, and stores it in the *Experience Collection File*. Periodically (typically daily) this information is consolidated with data from the previous 14 runs. This consolidated information is stored in the *Experience File* and is used to manage the workload in subsequent runs.

PCS collects experience information on every application known to the AMD. This allows you to build up experience information before deploying the other PCS features.

Battle Plan

A *battle plan* consists of the set of jobs to be run for any one instance of a production cycle², say on the first Monday of June of this year. It contains information about their scheduling details (from CA 7), PCS details (from the AMD), and history (from the Experience File).

Note: If the header job of a tree is submitted during the production cycle, then all of jobs in the tree are deemed to be part of the same battle plan, regardless if some of the jobs are submitted beyond the production cycle time frame.

A battle plan consists of any job submitted by CA 7 that belongs to an application in manage or monitor mode (the distinction is explained later in the chapter). PCS attempts to match the

submitted jobs to those identified by the PCS Forecast Collection Utility for this period. If the jobs can't be matched, they are known as ad hoc jobs. If a job does not belong to an application in managed or monitored mode, then it follows its non-PCS default behavior.

A particular run of a job, path, or job tree is referred to as a *job instance*, *path instance*, or *job tree instance*. If a job tree repeats during the production cycle, or has a name that only differs by schedule id, PCS processes each job as a distinct job instance. (This guide refers to *job* rather than *job instance*, *path* rather than *path instance* and *(job) tree* rather than *(job) tree instance*, using the latter term only when necessary for clarity).

Battle plans can be created once the information is available from the PCS Forecast Collection Utility. A battle plan has a status of *ready* from the time it is created until it is activated. During this time, you can override the defaults for the applications, trees, paths and jobs in this battle plan.

A battle plan needs to be activated for a production cycle, and auto-activation is a popular option. (Auto-activation occurs at the start time of the associated production cycle; alternatively you can activate the battle plan manually in the PCS dialog.) PCS then processes the jobs CA 7 submits according to this *active* battle plan.

There can only be one active battle plan at a given time. A battle plan is active until the end of the production cycle, or until another battle plan is activated. Thereafter the original battle plan is known as *closed*. If a battle plan is not activated during its production cycle, it is said to have *expired*. (A battle plan can persist past the end of its production cycle if any jobs have not completed; however it would be more typical for a new production cycle to start with a new battle plan and any incomplete jobs from the original battle plan to be automatically rolled over into the new battle plan.)

If necessary, you can create a one-time production cycle and when you (manually) activate its associated battle plan, it closes the current battle plan and the new one takes over.

Installation List

Datacenters often have a set of jobs they keep an eye on, perhaps because they are important, prone to fail, dependent on external triggers, and so on. PCS gives the datacenter a mechanism to group these jobs into a named *installation list* that can be chosen for display. Popular names for this list are "Loved Ones", "Watchlist" and so on. You can have more than one installation list. Being on an installation list does not influence the automatic management of the associated jobs.

Alerts

PCS provides a facility to raise an alert when any of seventeen events occurs. Possible events include a JOB TREE IS LATE, JOB CANCELLED BY CA 7, JOB BECAME CULPRIT. Alerts can be issued as an email, WTO, or CA 7 message.

Production Service Groups Considerations

You need to be aware of any Production Service Groups (PS Group) in your environment. PS Groups are a feature of ThruPut Manager AE to provide an method to direct jobs to the Production Services queue for processing. Jobs operating under a PS Group are similar to PCS-processed jobs in that they use the PS queue, and its associated Service Classes. They also share some settings in the SLM Policy. They are dissimilar in that they are not part of the PCS optimizations. Look for them in the Policy Configuration Options in the SLM dialog. At runtime, they appear in the SLM dialog on the Production Service Group Job Display panel, with their associated PS Group name. Further information on PS Groups is found in *TM AE Concepts and Usage Guides*.

PCS Workload Management Process

This section explains how PCS is designed to work once the setup is complete. It assumes the environment is ready and available, i.e., CA 7, JES2, and ThruPut Manager started tasks are all operational; the AMD has daily production cycles specified; and the battle plan is auto-activated.

Here are the highlights of the process for a job belonging to an application in monitor mode. Where the process is different for a job belonging to an application in managed mode, this is noted.

- CA 7 database is updated for the day.
- The PCS Forecast Collection Utility is run based on the updated CA 7 schedule.
- The battle plan is created. PCS determines the end time for all paths and milestone jobs, and the duration of all the jobs. From that it determines the expected start time (that's where the CA 7 definition of predecessor jobs and the PCS experience data comes in). Lastly, it determines the initial critical path and the initial urgent path for each job tree.
- The battle plan is activated.
- CA 7 places jobs in the REQUEST queue to await submission.(When such a job is not known by the PCS Forecast Collection Utility, it is referred to as an ad hoc job and acts as a single-job tree).
- PCS is notified when the job has entered the REQUEST queue and later again when it is submitted to JES2.
- As is usual with ThruPut Manager, the job is analyzed and processed by the JAL processor when it is submitted. JAL processes the job adding any binds, limits, messages and so on. All SLM statements are ignored for jobs belonging to an application in managed mode.
- A job belonging to an application in monitor mode is not processed in the PS queue by PCS. Rather it reverts to its non-PCS default processing, which means it is processed by one of:
 - JES2 if the submitting job class is not selectable by ThruPut Manager. (The JECL option JECLEXEMPT also influences whether the job is analyzed or not.)
 - JAL-assigned job class if the submitting job class is selectable and JAL processing does not assign the job to a Service Group.
 - ThruPut Manager AE, in the General Services queue, if JAL processing assigns the job to a General Service Group;
 - ThruPut Manager AE, in the Production Services queue, if JAL processing assigns the job to a Production Service Group.
- A job belonging to an application in manage mode is "processed" by PCS in the PS queue. PCS subcontracts SLM to manage the job according to the PCS parameters, which it may be updated as the workload progresses. SLM may change the Service Class, give preference to a job with respect to resources, or change its position in the queue.
- PCS keeps track of whether a job is on an urgent path, whether it's going to be late, whether it's a culprit for making a downstream job late, and so on. It sends alerts if one of 17 workload events occurs.
- PCS updates its displays for this job, its path, tree and application, as well as any affected jobs. The PCS job display includes an integrated view of JES2, CA 7 and ThruPut Manager status and information. ThruPut Manager's UDF (User Display Facility) provides further runtime information on each job.
- As each job ends, PCS ...
 - Updates the status fields,

- Updates "expected" values with "actual" values,
 - Updates the Experience Collection File,
 - Updates the SMF Record Collection file, and
 - Generates the Job Summary Report as part of the job output, if the appropriate option is set.
- The production cycle ends and the battle plan is closed. Normally a new production cycle is starting with its battle plan. Any outstanding jobs are inherited by the new battle plan.
 - Later, the datacenter runs a utility to transfer the Experience Collection File information for that production cycle to the Experience File for aggregation with the information from previous cycles, readying the Experience File for use with the next battle plan.

Ideally, during this process Production Control operators simply sit back and let PCS manage the workload in the optimal manner to best complete the workload on time. However changes in circumstances, priorities, and so on do occur. You have full control to make any adjustments as needed. See [Runtime View](#) below for more details.

Runtime View

Once workload is processed by PCS, you can observe the progress of your workload using the PCS display dialog, starting with the PCS "dashboard": View/Edit Active Battle Plan. This panel provides the overall workload statistics for this battle plan and production cycle. As you can see this is the battle plan for TUESDAY, and the production cycle covers the period from 8 a.m. to midnight. It was auto-activated at 7:55 a.m. There is 12 hours and 30 minutes left in this production cycle.

For this snapshot, at 11:30 a.m., one path is late with one job identified as the *culprit*.

There are 5 applications, with their 28 job trees, 36 paths, and 78 job instances to be completed (To-Do). Three of those jobs were culprits. Three applications are already complete (Done).

From here you can get more information on applications, or on job instances. You start at the application level with aggregate information and drill down to the tree level, the path level, and ultimately the job level.

Alternatively, you might choose to display a list of job instances that includes all jobs in this battle plan, just the ones on an installation list, just the list of jobs customized for this battle plan, or the list of culprit jobs that haven't yet completed. You can "drill up" to see information for the path to which a job belongs, the tree to which that path belongs, and the application to which that tree belongs.

To address different audiences, the runtime information is offered in three levels of capability:

- View current battle plan without making changes.
- View and edit current battle plan.
- Make any AMD changes, including to current battle plan.

```

- Action                                     Glossary   What's New   Help
-----
----- PCS -----
View/Edit Active Battle Plan

Command ===>

Production Cycle Name: TUESDAY                Type: Day of Week (Tue)
Description:
Cycle Date/Time - From: Nov 22 08:00          To: Nov 22 23:59
Activated: Nov 22 07:55                       by: **Auto**
... from AMD, Volser: DOCT.CA72.AMD MVSS03
Next Production Cycle: *Not Defined*
Current Date/Time: Nov 22 11:29:07           Time Remaining in Cycle: 12:30:53

-Applications-  -Job Trees-  -Job Paths-  -Job Instances-
Total Culprits
Is Late:        0          0          1          1          1
Will be Late:   0          0          0          0          0
Was Late:       0          0          0          0          0

To-Do:         5          28         36         78         3
Done:          3          14         22         48
Total:         8          42         58         126
Adhoc: 0

For Detailed Battle Plan Information, Select one of the following:
1 List of All Applications
2 Job Instances...
3 List of Incomplete Culprit Job Instances

X Exit

```

The information for each job includes:

- JES2 status information
- CA 7 status information
- ThruPut Manager information
 - Production Importance
 - Target and Acceptable Times, lateness information
 - Critical Path indicator
 - Urgent Path information
 - Lateness ("will be late", "late" and "was late") information
 - Why this job is (not) a culprit, affected jobs and so on

As jobs complete, PCS updates information to reflect the progress of the battle plan. It updates the dashboard statistics, replaces the expected end times with actual end times, and refreshes the status fields. For instance the Critical Path indicator may switch from the initial path to another path as the workload progresses. PCS collects experience data for integration with previous job statistics. It raises alerts whenever a relevant event occurs.

In managed mode, you also notice other data changing, reflecting the SLM management of the job. PCS and SLM are automatically reacting as the workload and environment change, based on the values and settings in the AMD and SLM policy. UDF and the Job Summary Report includes the relevant PCS information.

Operator Control and Intervention

Generally no actions are required on your part. However when the unexpected happens, you may need to react. You can:

- Stop the flow of jobs in the PS queue, and restore the flow when appropriate, using the following commands:

```
SLM SELECT NOPRODUCTION
SLM SELECT PRODUCTION
```

In addition, SLM SELECT NONE and SLM SELECT ALL apply to work in the PS queue (and the GS queue).

For instance you could use these commands when the online systems experience a particularly busy period, and you can afford to defer batch for awhile.

- Override the defaults and options at different levels. For instance you could increase the production importance, or change the time after which the job tree will be considered late.
- Add alerts to highlight milestone events or help diagnose problems.
- Set the *monitor-mode only* flag to make all jobs act as if their applications were set to monitor mode. PCS resumes normal processing of jobs in manage mode when the flag is reset.

Of course the existing ThruPut Manager interventions such as activating or deactivating binding and limiting agents still apply.

Implementation Considerations

You may implement PCS in a number of ways depending on

- Your current use of ThruPut Manager,
- How the target workload is running now,
- Your corporate conventions and facilities with respect to change management.

For instance, a typical scenario starts off with your workload already managed by ThruPut Manager AE and SLM. You have created General Service Groups, and possibly Production Service Groups, and your JAL assigns jobs to the appropriate group. You are moving to PCS to realize the benefits of managing (CA 7) scheduled jobs in the context of their schedules, relationships to other jobs and past experience.

You set up a few applications in monitor mode, leaving the rest "neither monitored nor managed". You remind everyone that they can see the PCS-processed jobs in the `PCS View Active Battle Plan` dialog. This way staff can see the PCS displays and information, as well as receive any alerts if a relevant event occurs. (Applications that are neither monitored nor managed don't appear in PCS displays, nor are PCS alerts triggered for them.) Jobs continue to be processed by their non-PCS default.

When you are comfortable with the PCS displays you can either change the status of all the applications to manage mode or a combination of manage and monitor mode. An application in manage mode means PCS takes actions to ensure the jobs that are running late get preference so their path and tree don't also become late. The Job Summary Report has PCS information included in it.

Eventually all the workload within the JESplex should be processed in either the General Services queue or the Production Services queue. When ThruPut Manager isn't managing all of your workload, you may get surprising results. For instance, since SLM holds back workload so more important jobs can get resources, a WLM job may be more aggressive and "take up" the apparently available capacity. When jobs are managed by ThruPut Manager, they all play by the same rules. Therefore we recommend that all jobs within a JESplex be managed by ThruPut Manager for optimum results.

During any implementation, consider your JAL:

- If PCS is down, jobs revert to their non-PCS default behavior. Your JAL should allow for this and assign all jobs to an appropriate service group.
- On the other hand, if PCS is active and a job belonging to an application in manage mode is submitted by CA 7, then all the SLM statements in your JAL are ignored for that job.

- Routing by class? If you have been relying on the job class to determine job routing, add a binding agent to accomplish this. (Refer to the *TM AE Usage Guide* for a more detailed explanation.)
- In all cases you should check your JAL for logic flow in light of any new workload.

Typical PCS Day

Once PCS is implemented and PCS is processing the applications you want it to, you settle into a new normal. Production cycles cover your batch hours and battle plans are auto-activated. From time to time applications are added to the CA 7 database and, in the PCS dialog, you refresh the applications, set the production importance and override any defaults for the new application.

Each day, once the CA 7 updates are completed, you run the PCS Forecast Collection Utility and create the ready battle plan. Any one-time changes to the ready battle plan are made. For instance you could add an alert, or increase the production importance for a specific application. This ensures the workload runs with up-to-date information. Whether there are changes or not, the battle plan is auto-activated.

During each production cycle PCS captures experience data in the Experience Collection File. A utility transfers the data from the Experience Collection File to the Experience File, and (re-)calculates the experience values for each job to be used for subsequent runs of the job.

Typically all the scheduled work is completed within the production cycle and the battle plan is closed. When jobs may extend past the production cycle they were created in, the battle plan persists until all the jobs are complete or until a new battle plan is activated. This new battle plan takes over the management of any lingering jobs.

Conclusion

PCS delivers superior proactive management for your CA 7 scheduled workload. It collects and infers information for applications, trees, paths and jobs. It uses the Application Management Database, production cycle and battle plan constructs to collect, organize and deploy this information. PCS subcontracts the actual processing of a job in managed mode to SLM, based on PCS values and settings.

PCS is a powerful product that has implementation steps and phases that need to be considered, though no step is particularly onerous or time consuming.

Chapter 1 of this manual deals with the PCS core features - any person implementing PCS needs to know this material, and this material covers the features needed to have any CA 7 scheduled workload managed by PCS. Further optional features of PCS are covered in [PCS Advanced Highlights](#).

Chapter 2

PCS Advanced Highlights

This chapter highlights some of the advanced features of PCS.

Prerequisites

The many advanced features of PCS are available to users with particular requirements or whose workload characteristics would benefit from further tailoring. They are mostly independent of one another, though they all build on the core implementation you have already realized.

Some topics further the behavior to your workload. These are covered below in [Advanced Highlights](#).

The rest of the advanced topics refine your implementation and these are covered, with appropriate context, in the *PCS Usage Guide*.

Advanced Highlights

Automated Capacity Management

If your datacenter has sub-capacity pricing arrangements with IBM or participating independent software vendors, Automated Capacity Management (ACM) controls batch load at peak times, based on the rolling 4-hour average (R4HA) of CPU usage to allow your datacenter to significantly lower monthly software license fees.

With ACM enabled, ThruPut Manager AE is aware of capacity limits. They may be defined as part of your soft capping process as LPAR Defined Capacity, LPAR Group capacity limits, or both. ACM also allows you to specify a capacity value for it to use as the capacity limit, either for a particular LPAR or set(s) of LPARs.

ACM monitors the R4HA of CPU usage and constrains or defers low importance batch work as the system approaches its capacity limit.

For more information on ACM in general, see *ACM Guide*; for PCS-specific information see the chapter in the *PCS Usage Guide*.

Quick Jobs

This feature may be of interest if you have very short running jobs that you would like to run "right away". It is appropriate to use when there is a mix of "quick" and "normal" jobs. If the environment has only short running jobs, the Quick Job treatment is unnecessary.

For jobs in the PS queue (belonging to an application in managed mode) you use the SLM Policy dialog to set the criteria to assign a job quick job treatment.

Quick jobs receive preferential selection processing: they are immediately aged to the to the C1 level, regardless of their production importance. Within the quick jobs at C1, they are queued by production importance, acceptable time and rerun factor (in that order).

Demand Job List

The PCS Forecast Collection Utility makes *scheduled* jobs available to a battle plan when it is created. You can tell PCS that a *non-scheduled* job tree might run by adding its header job to the Demand Job List. All jobs triggered by the header jobs directly or indirectly will be determined by PCS automatically.

Initially these jobs are seen as *unconfirmed* jobs in the battle plan. Once a header job is demanded in, they appear as *confirmed* jobs.

External Process List

An *External Process* is a PCS concept that recognizes that a CA 7 job is followed by process external to CA 7, which, when finished, triggers another CA 7 job. This triggering action is often based on the creation of a dataset by the external process and CA 7 using a dataset trigger to submit the downstream job.

Although CA 7 does not have any information linking together the job that initiated the outside work and the job triggered by CA 7 by the dataset, PCS can be instructed to build such a link. This can be done by creating an External Process.

Appendix A.

Production Service Groups

This appendix explains Production Service Groups from a PCS perspective.

Overview

PS Groups, from a ThruPut Manager AE perspective, are described in *TM AE Concepts and Usage Guides*.

Jobs assigned to PS Groups are processed in the PS queue along with any battle plans jobs belonging to an application in manage mode. Therefore they enjoy the same preferences over jobs in the General Services queue.

Jobs in PS Groups can serve as a stepping stone from SLM to PCS. If you plan to implement PCS after you have completed the SLM essentials, you may want to implement PS Groups for your production batch as part of the SLM implementation. This introduces the second batch queue and operations can get familiar with its behavior from the beginning.

PS Groups can also serve as a refinement once PCS is implemented. You may have production jobs from other schedulers that more properly belong in the PS queue rather than the GS queue.

Comparison with other jobs in the PS queue

Jobs assigned to a PS Group are similar to PCS jobs in the PS queue in that they:

- Use the PS queue, and its associated Service Classes.
- Have production importance specified for them.

They are unlike other jobs in the PS queue in that they:

- Interpret target and acceptable service times differently.

For a PCS job *target time* is the calculated *wall-clock time* a job should be completed by so that all job paths to which the job belongs remain at their target time or better. Ideally the target time of a path allows enough time to rerun the entire path before the acceptable time for the path is reached. When this cannot be achieved, the target time may be reduced and, in any case, it is never later than the acceptable time.

Acceptable time may be set for an application, tree or job, with PCS using these settings and other information to determine the *calculated acceptable time* suitable for the relevant jobs, paths and trees.

Target time and *acceptable time* for PS Groups are the *lengths of time* a job takes to age to the target and acceptable thresholds in their queue. Both these values are set by the installation.

- Are assigned to PS Groups in JAL (the PS Groups themselves are created in the SLM Policy).
- Do not appear in the PCS dialog.

