Oracle® Production Scheduling

Implementation Guide

Release 12.1

Part No. E05546-06

March 2011

ORACLE®

Oracle Production Scheduling Implementation Guide, Release 12.1

Part No. E05546-06

Primary Author: Gary O'Hara

Contributing Author: Paul Fenwick, Shane Hellsten, Alex Kim

Contents

Chapter 5

Chapter 9

Chapter 10

Manual Scheduling **211**

Chapter 11

Chapter 12

Chapter 13

Send Us Your Comments

Oracle Production Scheduling Implementation Guide, Release 12.1

Part No. E05546-06

Oracle welcomes customers' comments and suggestions on the quality and usefulness of this document. Your feedback is important, and helps us to best meet your needs as a user of our products. For example:

- Are the implementation steps correct and complete?
- Did you understand the context of the procedures?
- Did you find any errors in the information?
- Does the structure of the information help you with your tasks?
- Do you need different information or graphics? If so, where, and in what format?
- Are the examples correct? Do you need more examples?

If you find any errors or have any other suggestions for improvement, then please tell us your name, the name of the company who has licensed our products, the title and part number of the documentation and the chapter, section, and page number (if available).

Note: Before sending us your comments, you might like to check that you have the latest version of the document and if any concerns are already addressed. To do this, access the new Oracle E-Business Suite Release Online Documentation CD available on My Oracle Support and www.oracle.com. It contains the most current Documentation Library plus all documents revised or released recently.

Send your comments to us using the electronic mail address: appsdoc_us@oracle.com

Please give your name, address, electronic mail address, and telephone number (optional).

If you need assistance with Oracle software, then please contact your support representative or Oracle Support Services.

If you require training or instruction in using Oracle software, then please contact your Oracle local office and inquire about our Oracle University offerings. A list of Oracle offices is available on our Web site at www.oracle.com.

Preface

Intended Audience

Welcome to Release 12.1 of the *Oracle Production Scheduling Implementation Guide.*

See Related Information Sources on page viii for more Oracle E-Business Suite product information.

Deaf/Hard of Hearing Access to Oracle Support Services

To reach Oracle Support Services, use a telecommunications relay service (TRS) to call Oracle Support at 1.800.223.1711. An Oracle Support Services engineer will handle technical issues and provide customer support according to the Oracle service request process. Information about TRS is available at http://www.fcc.gov/cgb/consumerfacts/trs.html, and a list of phone numbers is available at http://www.fcc.gov/cgb/dro/trsphonebk.html.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible to all users, including users that are disabled. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at http://www.oracle.com/accessibility/.

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The

conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

Structure

1 Getting Started
2 Understanding Production Scheduling
3 Starting the Production Scheduling System
4 Configuring Desktop Components
5 Creating Production Scheduling Models
6 Setting Up Global Options
7 Configuring Scenario Properties
8 Configuring Solver Options
9 Evaluating Production Schedules
10 Manually Scheduling
11 Publishing Data
12 Integrating With Oracle Value Chain Planning Suite
13 Scheduling Maintenance Work Orders

Related Information Sources

Integration Repository

The Oracle Integration Repository is a compilation of information about the service endpoints exposed by the Oracle E-Business Suite of applications. It provides a complete catalog of Oracle E-Business Suite's business service interfaces. The tool lets users easily discover and deploy the appropriate business service interface for integration with any system, application, or business partner.

The Oracle Integration Repository is shipped as part of the E-Business Suite. As your instance is patched, the repository is automatically updated with content appropriate for the precise revisions of interfaces in your environment.

Oracle Advanced Supply Chain Planning Implementation and User's Guide

Oracle Advanced Supply Chain Planning provides information about integration between Oracle Production Scheduling and other Oracle applications.

Oracle Strategic Network Optimization Implementation Guide

This guide provides information on implementing and using Strategic Network Optimization.

Do Not Use Database Tools to Modify Oracle E-Business Suite Data

Oracle STRONGLY RECOMMENDS that you never use SQL*Plus, Oracle Data Browser, database triggers, or any other tool to modify Oracle E-Business Suite data unless otherwise instructed.

Oracle provides powerful tools you can use to create, store, change, retrieve, and maintain information in an Oracle database. But if you use Oracle tools such as SQL*Plus to modify Oracle E-Business Suite data, you risk destroying the integrity of your data and you lose the ability to audit changes to your data.

Because Oracle E-Business Suite tables are interrelated, any change you make using an Oracle E-Business Suite form can update many tables at once. But when you modify Oracle E-Business Suite data using anything other than Oracle E-Business Suite, you may change a row in one table without making corresponding changes in related tables. If your tables get out of synchronization with each other, you risk retrieving erroneous information and you risk unpredictable results throughout Oracle E-Business Suite.

When you use Oracle E-Business Suite to modify your data, Oracle E-Business Suite automatically checks that your changes are valid. Oracle E-Business Suite also keeps track of who changes information. If you enter information into database tables using database tools, you may store invalid information. You also lose the ability to track who has changed your information because SQL*Plus and other database tools do not keep a record of changes.

Chapter 1

Getting Started

This chapter discusses:

- Production Scheduling overview.

- Production Scheduling business processes.

- Production Scheduling implementation.

Production Scheduling Overview

Production Scheduling uses a constraint-based approach to automated scheduling. Unlike traditional automated scheduling tools that are limited to simple dispatch rules and have known bottlenecks, constraints in Production Scheduling can be assigned to every element-resource, operation, and due date in a schedule. Through advanced solver technology, feasible solutions can be found for virtually any constraint.

Production Scheduling is a finite scheduling product that addresses discrete batch processing with floating multistage problems. Production Scheduling can process an almost limitless number of variables and constraints. Using the Windows NT operating system, the total cost of ownership decreases.

Using the advanced scheduling capabilities of Production Scheduling you can:

- Create a model.

- Solve a model.

- Publish a production schedule.

- Dispatch schedule information.

Production Scheduling is designed to schedule discrete manufacturing environments that have the following characteristics:

- Many components comprise each product.

- Few final products exist in relation to the large number of components.

- Subassemblies can be out-sourced to third-party manufacturers.

- Many production stages exist.

- Multistage, floating production bottlenecks exist.

- Large numbers of customer orders exist, often make-to-order as opposed to make-to-stock.

- Products are produced in discrete batches or consistent lot sizes.

- Multiple operations and routings can exist for the production of an item, and the user is able to specify a preference.

- Minimum and maximum delays exist between work center operations.

- Setups, changeovers, and cleanups are required.

- Capacity constrained resources exist in production routings.

Unlike other scheduling systems, Production Scheduling has a self-configuring solver. This technology enables the solver to adapt as the model changes. It also includes item propagation, which immediately determines item feasibility; adaptive propagation, which starts and stops propagation only as it is needed; and effective backtracking.

Production Scheduling Business Processes

Through the seamless integration of Production Scheduling (PS) with Oracle Value Planning Suite (VCP), you can create high level supply chain plans and refine these plans to create optimal production schedules that can be implemented using Oracle transaction systems.

Using data from Oracle transaction systems and, if available, demand forecasts from Oracle Demand Planning, ASCP can create planned orders with a granular horizon and detailed multi-level pegging of supply and demand. PS can then take planned orders from ASCP and produce an optimized production schedule. You can then review and release the production schedule to Oracle transaction applications for implementation.

Using data from Oracle transaction systems and forecasts from Oracle Demand Planning, you can balance the conflicting objectives and limitations of supply, production, and distribution in your supply chain to determine how to meet demand with the least cost or with the most profit. You can also determine which facilities should be opened or closed, and in what order, throughout the horizon of a model.

Production Scheduling business processes are a part of the Oracle Value Chain Planning Plan to Produce business process. The Value Chain Planning Plan to Produce business process fulfills an organization's requirements for planning, deploying, producing, and assessing production. Production Scheduling is a part of the manufacturing execution phase of Plan to Produce. Production schedules produced by Production Scheduling are used as a basis for production deployment. To create a production schedule used to deploy production in an enterprise, this application relies on a series of interactive and non-interactive events.

This diagram illustrates the non-interactive and interactive business processes for Oracle Production Scheduling:
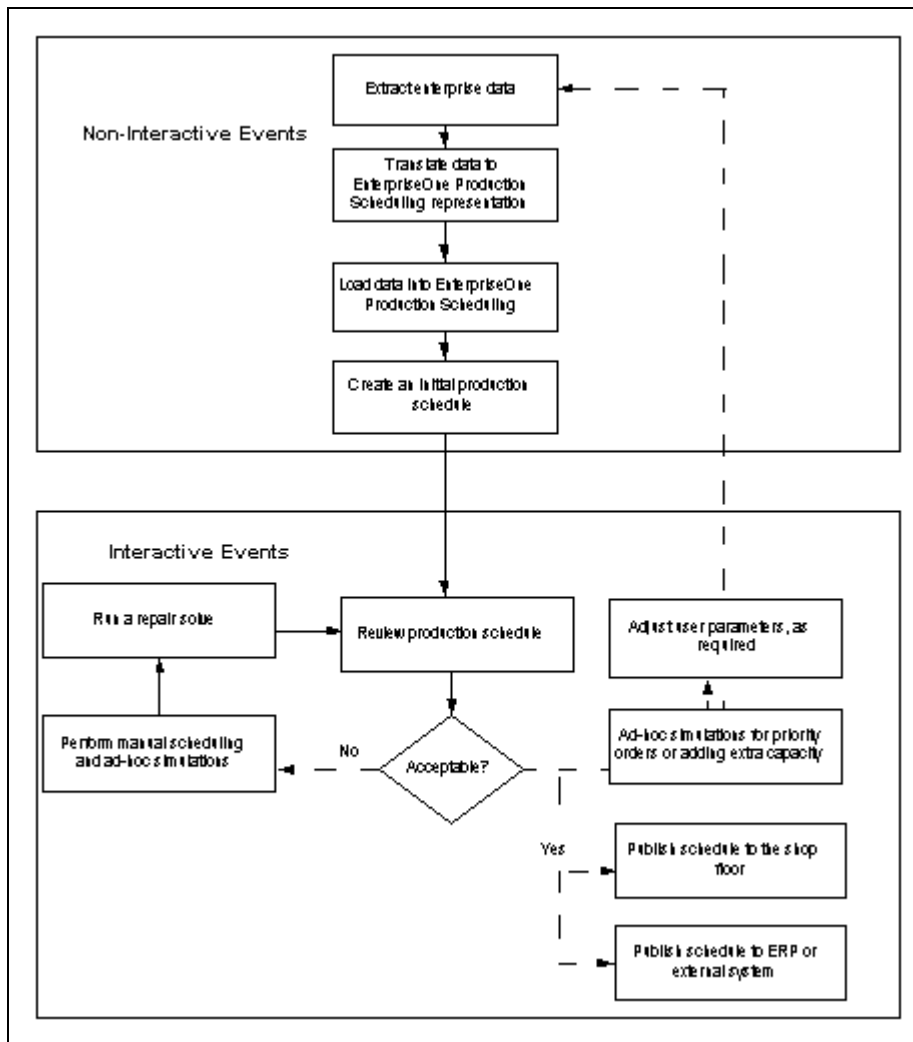
Non-interactive and interactive scheduling events

Non-Interactive Events

Production Scheduling business processes rely on non-interactive events to extract enterprise data from an ERP system and import it into the application. As a result of these non-interactive events, Production Scheduling is kept up to date with enterprise data in order to produce a production schedule that will be reviewed each morning to determine the daily production.

The following non-interactive events typically occur overnight or after the daily production shift has ended:

1. Enterprise data is automatically extracted from an ERP system.

2. Enterprise data is translated into an acceptable format for Production Scheduling.

3. Enterprise data is loaded into Production Scheduling.

4. The initial production schedule is produced and awaits review.

Interactive Events

After non-interactive events are complete for the scheduling cycle, interactive scheduling tasks can be performed to modify or streamline the schedule according to business needs. Interactive scheduling tasks begin with the initial production schedule that is created as a result of the non-interactive events that occur overnight.

1. Review the production schedule.

2. Decide if it is acceptable or not acceptable.

3. Perform manual scheduling and simulations that enable you to evaluate different scenarios (for unacceptable initial schedules).

4. Run a repair solve to incorporate the manual scheduling changes (for unacceptable initial schedules).

5. Perform simulations that represent the introduction of priority orders or extra production capacity into the schedule.

6. Publish the production schedule to the shop floor using the web dispatch list.

7. Publish the schedule to an ERP system.

Creating or Importing a Model

The creation of a production schedule requires a model of your production process to be either created or imported from an external system. The model represents the production capacity of your enterprise. A working model is the first step towards creating a feasible production schedule. It consists of the resources, operations, and routings that combine to form a representation of your production process.

An accurate production schedule requires up-to-date model data, including starting item levels, safety stocks, work orders, and supply and demand information. Model data is often updated on a daily basis.

The steps for building models, which are explained in more detail throughout the rest of this document, are only suggestions. You might find a different sequence that is better suited to your organization's needs.

Solving a Model

Solving the model produces a production schedule. You can make changes to the model and solve the model until a schedule that meets your requirements is produced. You can resolve scheduling problems by making changes to the data model including changing order priorities, increasing capacity, or by making manual changes to operations using what-if scenarios.

Additional information from the shop floor or other sources can also be reflected in the schedule. For example, real-time capacity changes, maintenance requirements, or supply and demand changes can exist.

Production Scheduling as a Service Call

Production Scheduling can be called as a service - the solver and application have been decoupled such that the import and solve process can optionally be run on the server – this includes Windows, Linux, Sun Solaris, HP-UX, and AIX platforms.

This feature enables Production Scheduling users to have their models pre-solved on the server – saving time and improving efficiency. It also allows for batch solving of multiple plans (for example, automation of the "what if" process).

The server-side functionality creates a deployable solved model which includes import and solver statistics in log files. The log files are all downloaded to the client desktop on planner request. This feature also includes improved application status logging and filtering - including access to various event statuses and solver statistics log from within Production Scheduling for events that took place on the server.

See Integrating Production Scheduling, *Oracle Advanced Supply Chain Planning Implementation and User's Guide*

Publishing a Production Schedule

Once you are satisfied with a production schedule, you can publish it for use on the shop floor. Once the schedule has been committed and published, Production Scheduling enables you to export operation schedules and reports to external systems.

Production Scheduling Implementation

The Production Scheduling implementation can be divided into the following phases:

• Install Production Scheduling.

• Define production data.

• Import enterprise data.

• Deploy production schedules.

Installing Production Scheduling

To install the Production Scheduling application, use the Production Scheduling Installation Guide.

Defining Production Data

To define production data you must:

1. Define resources.

2. Define operations.

3. Define routings.

See Creating Production Scheduling Models, *Production Scheduling*

Importing Enterprise Data

Data can be imported into Production Scheduling in various formats.

See Integrating With Oracle Value Chain Planning Suite, *Production Scheduling*

Chapter 2

Understanding Production Scheduling

This chapter discusses:

- Automatic floating bottleneck detection.

- Advanced analytical decision support.

- Direct scenario comparison.

Automatic Floating Bottleneck Detection

This section discusses:

- Production scheduling implementation

- Constraint-based scheduling limitations

- Constraint-directed search advantages

Production Scheduling Implementation

Production Scheduling reduces initial implementation and ongoing maintenance costs through a self-configuring solver. To understand this next-generation technology, you must consider current-generation approaches and their inherent limitations.

To implement current-generation scheduling products, you must adjust production schedules to make them usable. During this implementation adjustment phase, consultants hard-code rules for many aspects of the scheduling problem. By specifying bottlenecks, sequencing rules, and so on, you create a usable schedule. However, this approach is expensive and time consuming. Furthermore, the quality of the solution can degrade as the underlying assumptions of the hard-coded rules change. For example, changes to demand mix and capacity constraints are normal in any business, but they change the assumptions of the original rules and, subsequently, reduce the quality of the schedule. As schedule quality degrades, planners try to maintain quality by overriding the system with time consuming and subjective manual schedule changes. Planners often stop using the system entirely and revert to their previous planning methods.

Production Scheduling provides the ability to detect resource bottlenecks automatically, even when they dynamically float, or move, within a given schedule. Resource bottlenecks may float from crew to tool, and then to materials, to machine, and so on throughout the scheduling horizon. Using knowledge about resource bottlenecks, the Production Scheduling solver employs a scheduling strategy that exploits each type of bottleneck and maximizes manufacturing throughput. As a result, no initial rule configuration and schedule quality remains consistent over the life cycle of the implementation. Instead of focusing on customizing and maintaining rules, the planner can focus on scheduling.

Dynamic bottleneck detection is such an important concept in Production Scheduling that a specific view is dedicated to the visualization of bottlenecks while schedules are being solved. The Resource Contention view enables planners to watch the Production Scheduling solver dynamically detect and exploit resource bottlenecks. Red indicates serious resource criticality. Shades of orange, yellow, and green indicate decreasing criticality, with green indicating the least criticality. Once the solve finishes, the Resource Contention view displays a histogram of the resources that drove the majority of scheduling decisions.

For example, in this example, the planner can see that the red Electrical Tester machine is the dominant resource bottleneck at that point in the solve:

Resource Contention

The technology that enables users to detect bottlenecks automatically is called Constraint Directed Search (CDS). CDS is the logical evolution of constraint-based scheduling, but it does not have the limitations that historically affected constraint-based approaches.

Constraint-Based Scheduling Limitations

Constraint-based schedules provide an accurate and detailed model of the factory. Constraint propagation is also an efficient method for determining the consequences of scheduling decisions. However, although propagation usually finds a feasible solution, it often requires an excessive amount of time, particularly for large problems. These large problems are relatively common; therefore, the application of this technology in the real world has been slow because:

- A constraint-based scheduler must assign a start time to each operation and resolve resulting conflicts before a feasible solution can be reached.

  This delay causes the system to move data around rather than performing useful computation (known as *thrashing*) and leads to backtracking (undoing unworkable decisions), which significantly increases the time needed to find a solution.

- To prevent thrashing, implementers must customize algorithms for each implementation.

  Because this approach takes an existing set of rules and tries to fit those rules to a new problem, the algorithm requires numerous adjustments. These adjustments often require changes to the application source code, which compromises future upgrades and increases the time and cost of the implementation. In addition, custom algorithms cannot respond satisfactorily to automatic floating bottleneck detection.

- As a result of customization, current-generation technology is not generally applicable.

  If business requirements change, the technology cannot respond without further customization, delays, and expense.

- In many cases, current generation solvers cannot find good solutions even when they exist.

Constraint Directed Search Advantages

Constraint Directed Search (CDS) provides the advantages of constraint-based scheduling without the disadvantages of performance problems. This improved technology, which reduces guesswork and prevents wrong decisions, uses:

- Texture measurements to:

  - Determine which decision is most critical by measuring local schedule conditions, which identify bottlenecks before a decision is made.

  - Minimize or eliminate backtracking by returning to the last decision made or to the most uncertain decision that was previously made.

  - Allow larger problems to be solved.

  - Reduce solve time.

- Least commitment to:

  - Assign operation precedence constraints so that a start time does not need to be assigned.

    These constraints allow operations to float within a time range until the end, when an explicit start time is assigned.

  - Find the critical paths in a schedule.

  - Allow scheduling flexibility by tracking slack time.

  - Avoid deadlocks in the propagation process.

  - Optimize start time after sequencing.

  - Allow dynamic schedule monitoring, for example, to determine whether a late starting operation will cause the order to be late.

The Production Scheduling application supports both generative scheduling and repair-based scheduling. Production Scheduling allows forward and backward propagation as well as propagation in all dimensions of a scheduling problem, including resource utilization and inventory.

Advanced Analytical Decision Support

This section discusses:

- Production Pegging view.

- Resource Gantt and Resource Utilization combined view.

- Resource Gantt and Operation Gantt combined view.

- Resource Gantt, Operation Gantt, and Item Graph combined view.

Note. The views listed are only a subset of the analytical views that the Production Scheduling application supports. The eight basic views can be used on their own or combined through the use of a Wizard.

Production Pegging View

The Production Scheduling application provides many ways for planners to view their schedule. The Production Pegging view is used to analyze projected customer service levels and the supply constraints that dictate those levels.

All purchased and manufactured supply is pegged to a demand line item, which enables the planner to see which orders are late and what specific condition is preventing these orders from being on time. The Production Pegging view supports these key capabilities:

- User-defined demand sorting.

- Graphical-demand fill rates.

- Alerts with root cause.

- Dynamic-demand filtering.

User-Defined Demand Sorting

Demand sorting (for sales orders, forecasts, and so on) is defined by a user-defined folder structure, which is specified in the Supply & Demand Editor. The folder structure can represent product types, demand types, customer hierarchies, and so on, and it provides the planner with the most intuitive structure for analyzing the schedule.

Graphical-Demand Fill Rates

The folder summary bars have various colors that indicate the amount of on-time demand or the unit fill percent. Demands or demand folders with red summary bars indicate less than 33 percent of on-time units, orange summary bars indicate between 33 and 66 percent of on-time units, yellow summary bars indicate between 66 and 99 percent of on-time units, and black summary bars indicate 100 percent of on-time units. Position your cursor over a summary bar to see a description of the fill rate and details about late demand.

Alerts with Root Cause

Production Scheduling supports a variety of alerts. One type of alert warns the planner of projected late orders and their root causes. It also provides direct drilldown to the Production Pegging view, which automatically opens the Production Pegging view tree to the bottleneck operation that is pegged to the late demand.

Dynamic-Demand Filtering

Production Scheduling enables the user to filter thousands of orders by user-specific criteria, which, in turn, enables planners to manage by exception. For example, if eight of 31 line items occur late in a schedule, then a planner can filter the lines that are less than one day late. By filtering, the planner can learn that only four line items are more than one day late. All other orders are temporarily filtered out.

Resource Gantt and Resource Utilization Combined View

This view enables the planner to see the details of scheduled operations on their machines, crews, and tools, while viewing resource utilization for the currently selected resource in a bar chart. This view displays utilization by shift, daily, weekly, or monthly buckets. By selecting a particular time bucket, the planner can view the breakdown of idle, delay, down, changeover, and run time.

The Resource Utilization view also displays utilization across groups of resources. This view is useful for seeing the aggregate utilization of like resources, such as a pool of machines or crew.

Resource Gantt and Operation Gantt Combined View

This view displays how operations are scheduled. All operations for the selected resource appear in the lower pane, sorted chronologically.

Automatic sorting occurs based on the context of the zoom at the moment the resource is selected. Therefore, when zooming in or out, the planner can reselect the resource to re-sort the operations.

Resource Gantt, Operation Gantt and Item Graph Combined View

Although the Production Pegging view is ideal for a Make-to-Order environment, the Resource Gantt, Operation Gantt and Item Graph combined view is ideal for a Make-to-Stock environment. Examples of Make-to-Stock environments are Consumer Packaged Goods, and Food and Beverage.

The Resource Gantt and Operation Gantt panes in this view enable the planner to view the sequence being run on each resource. In the previous example, the planner can see that Packer 1 has a weekly campaign cycle for which time lost to sequence-dependent changeovers has been minimized. When you select each packing operation, the Item Graph pane displays the produced and consumed inventory levels, and inventory levels are perfectly supported.

Direct Scenario Comparisons

Production Scheduling can compare different schedule scenarios. During one session, a planner can create an unlimited number of scenarios and then compare them using the Key Performance Indicators (KPI) view.

The system provides standard supply chain metrics to assess which scenario best meets your business objectives, and a planner can sort scenarios by the KPI that they are interested in. Conversely, most competitive products support only a single scenario in one session, which forces planners to repeatedly save the various scenarios and run multiple sessions simultaneously, making the ideal scenario difficult to determine. In addition to KPI comparisons, schedules can be directly compared to each other from a customer service perspective. This type of comparison allows planners to understand which scenario best meets their customer service objectives. This capability provides answers to these key questions:

• In this new scenario, which work orders or demands are now on time compared to my best scenario previously?

• In this new scenario, which work orders or demands are now late compared to my best scenario previously?

- In this new scenario, what is the effect on overall earliness or lateness?

After a planner decides which schedule they prefer, they approve it and then publish it to a resource planning application.

Chapter 3

Starting the Production Scheduling System

This chapter provides an overview of system setup and discusses how to:

- Start Production Scheduling.

- Create, open, and import a scenario.

Understanding System Startup

You can start the Production Scheduling application from a Microsoft Windows desktop or from a command prompt, depending on your implementation. Starting the system from the Microsoft Windows desktop enables you to take advantage of the Production Scheduling desktop to create and evaluate scheduling scenarios.

The Production Scheduling application can be started directly from the command prompt or from a batch routine that can be automatically run at a specific time. The command prompt provides a number of options that enable you to perform any or all of these functions:

- Start the Production Scheduling application.

- Create, name, and save a schedule.

- Write the command and solve details to a log file.

- Terminate the session.

Starting Production Scheduling

This section discusses how to:

- Start the system from the Microsoft Windows desktop.

- Start the system from a command line prompt.

- Check the version of the software.

Starting the System from the Microsoft Windows Desktop

To start the system from the Microsoft Windows desktop:

1. Access the Microsoft Windows desktop.

2. Click the Start button, and then go to Programs, Oracle Value Chain Planning *x.xx.*

3. Select Production Scheduling.

Starting the System from a Command Line Prompt

To start the system from a command line prompt:

1. Access the command line.

2. Navigate using this path: <PS install directory>\SCP\<version>\Common\Start.

3. Enter *PsLauncher.*

Checking the Version of the Software

Access the Production Scheduling desktop.

1. From the Help drop-down list box, select About Production Scheduling.

2. In the About Production Scheduling window, click Version Information.

Creating, Opening, and Importing a Scenario

This section explains how to:

• Create a new scenario.

• Open or import and export an existing scenario.

Creating a New Scenario

To create a new scenario:

1. Access the Production Scheduling desktop.

2. From the File menu, select New.

   A new scenario is created in the Model Workspace.

Opening or Importing and Exporting an Existing Scenario

This section discusses how to:

- Open or import a scenario.

- Import a scenario.

- Export a scenario.

Opening or Importing a Scenario

To open or import a scenario:

1. Access the Production Scheduling desktop.

2. From the File menu, select Open and browse to your directory.

3. From the list of scenarios, select the one you want to open and click Open.

   Alternatively, if no scenarios are listed, browse to the location where you saved the scenario. The scenario can be an .xml, .xml.gz, .ops, .opz, or .dxt file. A .dxt file is a collection of scenarios. Files with the .ops or .opz extension are equivalent to xml and xml.gz files, respectively.

4. If you have a scenario open and want to open another, the system will ask if you want to import the file into the existing model.

   Select one of these options:

   - Click *Yes* if you want the first scenario to remain open.

   - Click *No* if you want to close the first scenario. You are prompted to save the scenario.

If the scenario you are opening or importing is an XML file, the scenario opens along with the associated Gantt charts. If the file is a DXT file, the Gantt chart does not open until a scenario is opened.

Importing a Scenario

There are two options to importing a scenario. You can follow the instructions detailed in the Opening or Importing a Scenario section, or you can:

1. Right-click on a scenario folder.

2. Select Import Scenario.

Exporting a Scenario

To export a scenario:

1. Right-click on the scenario.

2. Select Export Scenario.

Sending Scenario via E-mail

This topic provides information on sending a scenario file from Production Scheduling.

Sending Scenario via E-mail

To send a scenario file via e-mail directly from Production Scheduling:

1. From the Model Workspace, right-click a scenario and select Send Scenario to Mail Recipient.

    Your e-mail application opens with the scenario file attached to it.

2. Select the appropriate e-mail recipients and send the e-mail.

3. To open the scenario file, double-click the e-mail attachment.

    The recipient must have the same version of production scheduling installed to view the file.

Chapter 4

Configuring Desktop Components

This chapter provides an overview of the Production Scheduling desktop and discusses how to:

*   Configure the model workspace.

*   Configure toolbars.

*   Configure schedule views.

Understanding the Production Scheduling Desktop

This section discusses:

*   Data editors.

*   Key performance indicators (KPI) view.

*   Workbook.

*   Log window.

*   System icons.

The Production Scheduling desktop is the main user interface for creating production schedules and solving production schedule problems.

The Production Scheduling desktop is the principal user interface for all production scheduling activities. The desktop interface enables you to create and revise production schedules, solve production schedule problems, and view graphical representations of schedule data. This example shows the principal areas of the Production Scheduling desktop:
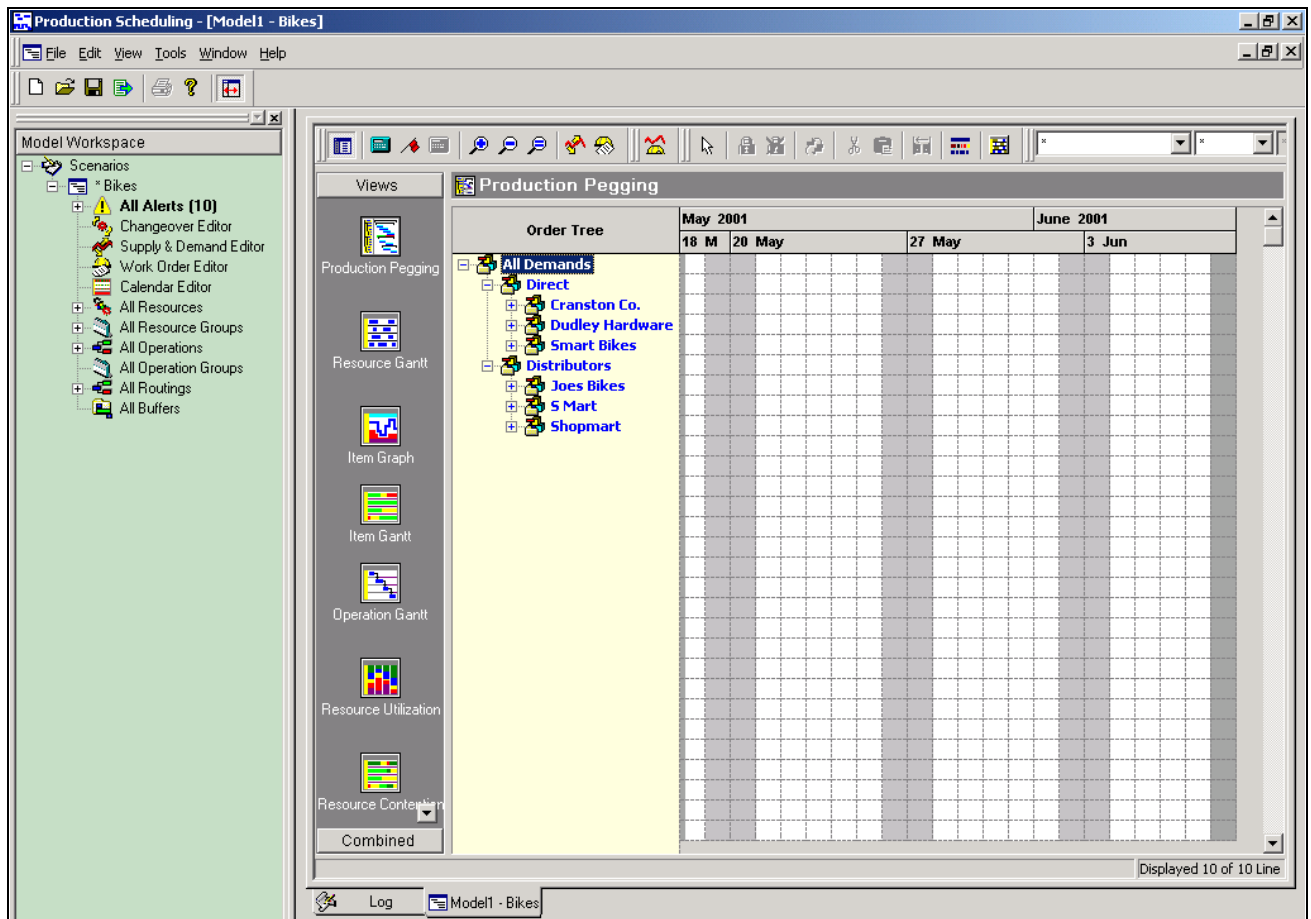
Production Scheduling desktop

The File, Edit, View, Tools, Window, and Help menus appear at the top. Immediately below these menus are the Production Scheduling buttons that you use to perform a variety of actions when you are working with your schedule. Below the buttons is the Model Workspace, which contains the tree structure of your scenario.

The scenario workspace tabs appear at the bottom of the screen. They display which work spaces are open. The Message bar is in the bottom right corner.

The majority of the screen is occupied by the views of your scenario, such as the Gantt chart, Item graph, and so on.

Data Editors

Data editors are views that enable you to record events, administer orders, and monitor supply and demand information. The data editors are:

- Changeover Editor.

- Supply & Demand Editor.

- Work Order Editor.

- Calendar Editor.

Changeover Editor

The Changeover editor enables you to record precise information about the time required to modify a machine between different operations. Using the Changeover editor, you can define relationships that enable you to define sequence dependent changeover relationships. In addition, you can create groups to save time when entering changeover information. All of the information entered into the Changeover editor is automatically saved in the schedule.

Supply & Demand Editor

The Supply & Demand editor enables you to add, delete, change, or view the properties of any supply or demand order or line item. It also enables you to add details about customers. Supply and demand information is usually imported.

Work Order Editor

The Work Order editor enables you to add, change, delete, and view production orders, maintenance orders, and engineering change orders. Using the Work Order editor, you can create, delete, and modify work order views.

Calendar Editor

The Calendar editor enables you to record downtime events and delay time events. An operation must end before a downtime event starts. An operation may be interrupted by a delay event and will resume after the delay event has completed.

Resource availability for crews, machines, tools, and machine changeovers runs according to a specified calendar, which enables you to create as many calendars as necessary to accurately communicate the time constraints that affect your resources. No downtime or delays are scheduled in the default calendar that ships with Production Scheduling.

The Calendar editor enables you to create recurring patterns, select multiple events in a user-defined range, and edit all events at one time.

Key Performance Indicators (KPIs) View

The KPI view displays schedule performance measurements in the Workbook area. Performance indicators for all schedules in the model can appear at the same time or be analyzed to view a particular schedule. KPIs are organized by the categories Customer Service, Materials, Costs, and Manufacturing metrics.

By adjusting scenarios and then comparing KPIs, you can determine the effect of the adjustments on your schedules. In addition, you can compare a baseline schedule to another schedule using the comparison capabilities at the bottom of the KPI view.

See Also

"Working with Key Performance Indicators," Evaluating Production Scheduling, *Production Scheduling*

Workbook

The Workbook area displays object windows, model editors, schedule views, and system messages. The Workbook area provides flexibility for viewing different types of Production Scheduling information.

You can display the Workbook area with or without tabs for each open diagram or model editor. On the View menu, click Workbook to open as many as 26 windows simultaneously and to display the Workbook tabs, which enable you to quickly identify and open the different views.

Each view can have several tabs. For example, you can have several schedule views open at the same time, as well as several Model editors open at the same time. The tabs automatically resize to fit your screen. You can hide the Workbook tabs by clearing the Workbook option from the View menu.

With Workbook tabs hidden, you can still toggle between windows by making a selection from the Windows menu.

Log Window

A Log window launches by default every time the system is launched. The Log window displays output from the solver file at the end of the solve. This output consists of a report about the solve and any problem messages. Right-click in the Log window to navigate to the log files.

System Icons

System icons are compact graphical elements that denote items such as activities, groups, folders, and resources in the Model Workspace. These icons help you quickly identify objects in your model. This table describes each system icon:

Icon Name	Description
Operation or Routing folders	Represents either an operation or a routing folder. Both operations and routings can be organized in folders for better accessibility.
Operation	Represents an operation to which resources are attached. The icon appears automatically in an Object window when the Open Diagram menu option is selected. This icon can be dragged from the Model Workspace into a routing diagram.
Operation Group folder	Represents an operation group folder. Operation groups can be organized into operation group folders for better accessibility.
Operation or Resource Group	Represents resources that are grouped for analysis purposes. Machine resources groups can be available within the Changeover editor. Like resources, operations can be grouped together in operation groups.

Icon Name	*Description*
Routing	Represents operations and routings that are connected in a specific order.
Operations, Resource, or Visits folder	Represents an organization, resource, or visits folder. Operations, resources, and visits can be organized in folders for better accessibility.
Visit	Represents a group of aircraft maintenance requirements and tasks assigned to a given organization. Visits are typically comprised of logical milestones, which are groups of tasks that must be finished before proceeding with the next milestone, or groups of tasks.
Crew	Displays details about Crew resources, including properties. You can drag Crew resources into operation diagrams from the Model Workspace.
Item	Displays details about Item resources, including properties. You can drag Item resources into operation diagrams from the Model Workspace.
Storage Space	Displays details of Storage spaces, including the name of the Storage space and its properties.
Machine	Displays details of Machine resources, including properties. You can drag Machine resources into operation diagrams from the Model Workspace. Machine resources appear in the Changeover editor.
Vendor	Displays details about Vendor resources, including properties. You can drag Vendor resources into operation diagrams from the Model Workspace.
Tool	Displays details about Tool resources, including properties. You can drag Tool resources into operation diagrams from the Model Workspace.
Schedule folder	Represents schedules that are organized into folders.

Configuring the Model Workspace

This section provides an overview of the model workspace and discusses how to:

- Hide the model workspace.

- Customize the model workspace.

- Activate workbook menus.

- View multiple diagrams.

Understanding the Model Workspace

The model workspace contains a tree structure, which contains all the scenarios that make up your model, and the operations, resources, and schedule views that make up each scenario.

When you create a new scenario, you need to define all of your resources first, then all of your operations, and finally all of the routings in your model. You can create several different scenarios to optimize your schedule. Create what-if scenarios that duplicate the scenario, and then modify the schedule and see how the modifications affect production. Once you have a satisfactory production schedule, save the modifications as a new model. To increase the available workspace in the Workbook area, hide the scenario workspace.

The model workspace contains one or more scenario folders for your model. Each scenario folder contains all of the individual scenarios and schedules that you have created for the current model. Scenarios allow the operations and resources of a model to be manipulated to create what-if situations. You can analyze different scenarios using various KPIs to determine the ideal schedule.

When you create a new scenario, Production Scheduling automatically sets up these sub-folders to help you organize the information:

- All Alerts

    See "Working with Alerts," Evaluating Production Schedules, *Production Scheduling*

- All Resources

    See "Defining Resources," Creating Production Scheduling Models, *Production Scheduling*

- All Resource Groups

    See "Defining Resources", Creating Production Scheduling Models, *Production Scheduling*

- All Visits

    See "Defining Visits," Creating Production Scheduling Models, *Production Scheduling*

- All Operations

    See "Defining Operations," Creating Production Scheduling Models, *Production Scheduling*

- All Operation Groups

    See "Defining Operations," Creating Production Scheduling Models, *Production Scheduling*

- All Routings

    See "Defining Routings," Creating Production Scheduling Models, *Production Scheduling*

- All Buffers

    See "Defining Routings," Creating Production Scheduling Models, *Production Scheduling*

In addition to the folders, the scenario contains these data editors:

- Changeover editor

    See "Defining Changeover Rules," Creating Production Scheduling Models,*Production Scheduling*

- Supply & Demand editor

  See "Specifying Supply and Demand Events," Creating Production Scheduling Models, *Production Scheduling*

- Work Order editor

  See "Working with Work Orders," Creating Production Scheduling Models, *Production Scheduling*

- Calendar editor

  See "Defining Work Patterns and Availability (Resource Calendars)," Creating Production Scheduling Models, *Production Scheduling*

Schedule Groups

Schedule groups enable multiple operations and resources to be combined as a group when viewed in the model workspace and in the schedule views. Using the model workspace, you can organize resources and operations into logical groupings to help you organize your model data. You can create a default resource group for each schedule that appears when views are first opened.

All schedule views except the Production Pegging view have a drop-down menu that lists the defined groups. When you select a group from the menu, the schedule view changes to display only those resources or items belonging to the group. Each schedule view can have its own default group view. The current default group in each view appears with a check mark next to its name.

Hiding the Model Workspace

To hide the model workspace:

1. On the menu bar, click View.

   Model Workspace is selected by default.

2. Click Model Workspace.

   The Model Workspace window is hidden.

3. Repeat steps 1 and 2 to alternately display and hide the Model Workspace.

Customizing the Model Workspace

To customize the Model Workspace:

1. Access the Model Workspace.

2. Right-click any blank area in the Model Workspace.

3. From the menu, select an option.

Allow Docking	Select to allow the Model Workspace window to move anywhere on the screen and attach to any side. To dock the model workspace to a specific side, move the window as close as possible to that side. The window expands to full size and locks in place when you release the mouse.
	When Allow Docking is cleared and Float in Main Window is selected, the model workspace moves anywhere within the Production Scheduling system window and does not attach to a side.
	When this option is not selected, select the Float in Main Window option to ensure that the right-click menu is available in the model workspace.
	Clear the Float in Main Window option and select the Allow Docking option when you want to dock the Model Workspace window again.
Hide	Click the window's close button to hide the model workspace. To display the model workspace again, select Model Workspace from the View menu.
Default Background Color	Click to return the background window color to the default color.
Background Color	Displays a color palette from which you can select a new background color. You can also define a custom color, which you can save.
Float on Top	Select to move, or float, the model workspace anywhere on your screen, including beyond the edge of the Production Scheduling system window. If the Allow Docking option is selected, then the model workspace attaches (docks) to any side of the screen. If the Allow Docking option is cleared, then the model workspace does not dock to any side of the screen. When this option is selected, the right-click menu is not accessible in the model workspace.
Float in Main Window	Select to move, or float, the model workspace anywhere inside the system window.

Activating Workbook Menus

To activate workbook menus:

1. Access the Workbook area.

2. Open two or more diagrams or editors.

3. From the View menu, select Workbook.

   The Workbook tab disappears.

4. From the Window menu, select one of the listed views.

Viewing Multiple Diagrams

To view multiple diagrams:

1. Access the Model Workspace.

2. Click the operations routings to open two or more diagrams.

3. Select Windows from the Window menu.

   The Window List appears, displaying all currently open windows.

4. Press Ctrl and select the windows that you want to view.

5. Select one of these options:

   • *Tile Vertically:* Select to display the diagrams vertically, that is, tiled vertically in the same window.

   • *Tile Horizontally:* Select to display the diagrams side-by-side.

6. Click OK.

Configuring Toolbars

This section provides an overview of toolbars and discusses how to:

• Hide toolbar options.

• Create a custom toolbar.

Understanding Toolbars

Toolbars contain buttons that represent each tool that you use in Production Scheduling. To understand what each tool does, pass your cursor over each active or inactive button. Descriptive help appears in the tooltips and the status bar.

Production Scheduling enables you to build a model of your process. By using the Production Scheduling tools located in the toolbar, you can create a graphical representation of the model. After the model is built, you can then access and view every element of the model with just one mouse click.

The Production Scheduling toolbars are context-sensitive; therefore, different sets of toolbars appear based on the task that you are performing. For example, the Alignment toolbar appears only when you are working with an Operation diagram.

The Customize window enables you to select toolbar options. You can use these options to select or clear tooltips, to select a display style and size for buttons, and to create or delete toolbars. Using the Customize window, you can:

• Select Show Tooltips to view help about the tool that is currently selected.

• Select Cool Look to remove the button-like appearance.

• Select Large Buttons to display enlarged buttons in the toolbar.

Tools are arranged in sets on the toolbar. Use different sets when navigating different views. This table describes the toolbars that are available in the Production Scheduling desktop:

Toolbar Name	*Description*
File	Use this toolbar to: • Create a new model. • Open or save a model. • Print an active document. • View help files. • Hide or view the Model Workspace.
Zooming/Planning	Use this toolbar to: • Zoom in. • Zoom to fit all components. • Fit the current selection. • Move the entire drawing area.
Resource	Use this toolbar to: • Select or move a diagram component. • Insert text or an image file. • Link multiple diagram components. • Insert an item set or a resource set. • Arrange components from top to bottom or from bottom to top. • Arrange components from left to right or from right to left.
Alignment	Use this toolbar to: • Align the top of two or more operations. • Horizontally center two or more operations. • Align the bottom or two or more operations. • Left-align two or more operations. • Right-align two or more operations. • Vertically center two or more operations.

Toolbar Name	Description
Nudging	Use this toolbar to: • Move one or more objects up one point. • Move one or more objects down one point. • Move one or more objects left one point. • Move one or more objects right one point.
Structure	Use this toolbar to: • Move an object to the front. • Move an object to the back. • Bring an object forward. • Move an object backward.
Supply and Demand	Use this toolbar to: • Show or hide the supply and demand tree. • Copy an order. • Paste an order. • Duplicate an order. • Show or hide pegging information. • Refresh pegging information.
Key Performance Indicator	Use this toolbar to customize KPI colors.

Hiding Toolbar Options

To hide toolbar options:

1. Access the Customize window, and click the Toolbars tab.

2. Select a toolbar from the list, and then select the toolbar options that you want to use on the toolbar.

   Toolbar options that are cleared do not appear on the toolbar.

3. Click OK.

Creating a Custom Toolbar

To create a custom toolbar:

1. Access the Customize window, and click the Toolbars tab.

2. Click New.

3. Specify a name for the toolbar and click OK.

4. Select the Commands tab and then select a category.

5. Select a button and drag it to the new toolbar.

6. Click OK.

Configuring Schedule Views

This section discusses how to:

*   Create new group views.

*   Configure the Production Pegging layout.

*   Set Multi-Capacity Resource Graph colors.

*   Set Item Graph colors.

*   Set Item Gantt colors.

*   Set Resource Contention view colors.

Creating New Group Views

To create a new group view:

1. Access a Schedule window.

2. Right-click the List Bar and select Add New Group.

3. Enter a name for the new group and click OK.

    The new group is added to the List Bar window.

4. Click the new group in the List Bar window to display its affiliated views or to add new views.

Configuring the Production Pegging Layout

To configure the Production Pegging layout:

1. Access the Production Pegging view.

2. Right-click the Order Tree header and select Layout Options.

3. In the Production Pegging Layout dialog box, select a layout type:

    • Show individual routing instances.

      Each operation and routing instance appears in a single line.

    • Show merged routing instances.

      All routings of the same type are grouped in a single line.

4. If you selected the layout type for merged routing instances, select one or both of these options:

    • Group multiple operation lines in a folder.

      All operations of the same type are grouped together in an operation folder.

    • Automatically expand operation folders.

      Operation folders automatically expand to display their contents. This option is available only if you selected the option to group multiple operation lines in a folder.

5. Show resource alternates as a set.

Setting Multi-Capacity Resource Graph Colors

To set Multi-Capacity Resource Graph colors:

1. Access the Edit Colors window.

2. Double-click one of these options to configure the color scheme:

    • Graph Background

    • Used Capacity Graph

    • Maximum Capacity Graph

3. Click OK to change the color scheme.

Setting Item Graph Colors

To set Item Graph colors:

1. Access the Color Selection window.

2. Select an appropriate color option from the available options:

    • Graph Color

    • Minimum Level Color

    • Maximum Level Color

    • Background

3. Click the Close button in the Color Selection window to exit.

Setting Item Gantt Colors

To set Item Gantt colors:

1. Access the Customize Colors window.

2. Click and drag a marker to the desired percentage.

   By default, the system displays color representations of inventory levels between −10 and 110 percent of the minimum and maximum levels.

3. To assign a color to a percentage range, right-click the color bar.

4. In the Color dialog box, select a new color and click OK.

5. Click Update Slider when you have completed the changes to the color scheme.

6. Click Apply and then click OK.

   Note. To return all of the values to the default colors and to the default number of color markers, click the Default button.

Setting Resource Contention View Colors

To set the Resource Contention colors:

1. Access the Customize Colors window.

2. Set these values:

   • In the Min. field, enter a value for the minimum percentage level.

   • In the Max. field, enter a value for the maximum percentage level.

3. To add markers, enter a new value in the # Markers field, and then click the color bar in which you want the new markers to appear.

4. To remove markers, enter a new value in the # Markers field.

   The appropriate number of markers is removed from the color bar from right to left.

5. To assign a color to a percentage range:

   • Right-click the color bar.

   • In the Color dialog box, select a new color, and then click OK.

6. Click Update Slider when you have finished making changes to the color scheme.

7. Click Apply and then click OK.

Note. To return all of the values to the default colors and to the default number of color markers, click the Default button.

Chapter 5

Creating Production Scheduling Models

This chapter discusses how to:

- Define resources.

- Define operations.

- Define routings.

- Define work patterns and availability (resource calendars).

- Specify supply and demand events.

- Work with work orders.

- Define changeover rules.

Defining Resources

This section provides an overview of resources and discusses how to:

- Define attributes for all resources.

- Define resource general properties.

- Define resource capacity.

- Define throughput.

- Define resource solver options.

- Define resource availability.

- Define resource costs.

- Define resource attributes.

- Define item constraints.

- Define item costs.

- Define item attributes.

- Define item suppliers.

• Define resource groups.

Understanding Resources

This section discusses:

• Resources.

• Batchable resources.

Resources

The first step in creating a model involves setting up the resources to use. Before attempting to create a schedule in Production Scheduling, all of the relevant resources must be entered in the system. Resources are any person, place, or thing that is used to perform an operation. Typically, only constraining resources are setup within your model, non constraining resources may be modeled optionally for visibility purposes.

Production Scheduling enables you to organize resources using resource folders. Resource folders organize your resources into logical groupings for quick access. For example, you can create folders for crews, items, machines, storage spaces, tools, and suppliers. To keep your model organized, create relevant folders to store all of the resources of a particular type. The resource folders can also be deleted if they are obsolete.

Resources are created in the Resources folder within the Model Workspace. Detailed information about each resource is collected and recorded in its properties, including its capacity, availability, cost, constraints, attached documents, and the operations and work orders where it is used.

You can also set up resource groups. Resource groups let you specify two or more resources that are represented together in schedule views.

After you enter resources, you have to specify resource properties including capacity, availability, costs, and constraints. Depending on the type of resource that you are configuring, the type of details required can vary.

The standard resource categories are:

Crews	A crew is an individual or group of people who work together to perform an operation. This individual or group works according to a specific calendar that reflects working hours, down time, or vacation time. Any number of calendars can be maintained by Production Scheduling to ensure that the resource is accurately represented for scheduling purposes.
	The properties for a crew resource enable you to rename the crew, add notes, and set capacity, availability, define solver options, and costs.

Items	Items can be both produced and, or, consumed in a given manufacturing operation.
	Production Scheduling records a great deal of information about each item used by the system, including inventory levels, storage locations, and any costs associated with the part. Each item can further be identified as either manufactured or purchased, flagged as saleable, and associated with various attributes for operation batching. You can associate purchased items with suppliers and indicate supplier preferences.
	Storage spaces typically contain many items. You should define all necessary storage spaces prior to entering your items.
Machines	Machines that are required to perform an operation are defined as resources. Machines are affiliated with a calendar that defines when down time or delays are scheduled to occur.
Shared Storage Spaces	Shared storage spaces are locations where items can be stored. A storage space can contain more than one item.
Tools	Tools are anything required to complete an operation. Tools, such as machines, are available according to a calendar that specifies when down time or delays are expected.
	The properties for the resource enable you to rename the resource, add notes, and set capacity, availability, and costs.
Suppliers	Suppliers are defined as any vendor who supplies purchased items or materials that are required to complete an operation or routing. Supplier data includes lead times, order multiples, and supplier preference in the event that multiple suppliers are available to supply the same purchased item or material. When you define a purchased item in the model, you can associate the item with a supplier. However, suppliers must be defined in the model before they can be associated with purchased items.
	This category acts independently of the Supply and Demand editor where supply events are entered. It can provide a quick way to theoretically replenish inventory from a vendor at a strategic point in the schedule without specifying all of the details necessary for creating a purchase order.

Batchable Resources

A batchable resource can be described as a machine, oven, or kiln that processes many items simultaneously. Batchable resources must share one or more similar attribute such as temperature, color, or duration. Typically, these resources have several spots, such as locations or racks. When an item is placed within a batchable resource the item consumes one or more of these spots. Depending on the characteristics of the operation and resource, the door to the batchable resource can be periodically opened throughout the duration of the operation, or the door must remain closed until the operation is complete.

When defining a resource as *Batchable,* you must define a minimum and maximum spot capacity, and a pull forward window that allows a resource to be filled to its minimum spot capacity. If a minimum spot usage is required before setting up a machine to run a certain job type, the Production Scheduling solver may require to pull-forward subsequent operations to achieve the target minimum spot usage. However, it is important to note that target minimum spots are considered a soft constraint. This means that if no operations can be pulled forward to fulfill the target minimum usage (due to the pull forward window) or if there isn't enough demand in the model, the target can be violated and the user is presented with an alert.

See "Defining Resource Capacity," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

When defining item attributes, you associate one or more attributes and attribute values to the item. For example when processing aluminum, it is required to cook certain items together at the same temperature for the same duration to achieve certain alloy characteristics. Items that share the same attributes and values are grouped together, for example both temperature and duration. If you decided to group by one attribute type, the schedule may include several jobs running sequentially within a batchable window, which allows an oven door to open and close to put items in and out during the period in which the oven is heated to a certain temperature.

See "Defining Item Attributes," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

Operations that use a batchable resource are grouped by attribute. If this is one attribute (such as temperature), all operations that have the same temperature are grouped together up to the maximum capacity of the resource. If the user groups by more than one attribute (for example temperature and duration), all operations sharing the same duration and temperature are run together to the maximum capacity of the resource.

See "Defining Resource Attributes," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

Setting up batchable resource attributes is a milti-step process. This list provides the high-level steps to setting up batchable resources:

1. Set up the resource attribute, such as duration or temperature, using the Attribute Properties window.

   See "Defining Attributes for All Resources," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

2. Set up the item attributes using the Item Properties window.

   See "Defining Item Attributes," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

3. Set up one or more resources as a batchable resource, and determine the resource attributes using the Resource Properties window.

   This includes the Resource Properties - Capacity tab and the Resource Properties - Attributes tab.

   See "Defining Resource Attributes," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

4. Set up the operation.

    If more than one resource can run the operation, the user must define a resource set. When defining the resource set, on the Link Properties page, the Capacity Required field should indicate the number of spots consumed by an item. For example, if an item takes up two spots in an oven, the Capacity Required field would be set to 2.

    See "Creating Item and Resource Sets," Creating Production Scheduling Models, Defining Operations, *Production Scheduling*

Defining Attributes for All Resources

To define attributes for all resources:

1. From the Model Workspace, right-click on All Resources.

2. Select New Attribute.

3. Enter the name of the attribute, such as Temperature or Duration.

    The system places attributes in alphanumerical order.

4. Right-click on the attribute to access the attribute Properties window.

5. From the attribute Properties - General tab, enter these fields:

Code	Enter the attribute code.
Notes	Enter relevant notes.

6. Select the Values tab and enter these fields:

Attribute Value	Enter all valid attribute values. For example, if the resource attribute is Temperature, attribute values could be 500 degrees, 550 degrees, 600 degrees, and so on.
	Attribute values appear below the attribute within the tree of the Model Workspace.
	Note. You can define as many attribute values as needed. Attribute values can be deleted or renamed.
Notes	Enter relevant notes.

7. Select the Where Used tab:

Items	Displays all items that the attribute is associated to.
Resources	Displays all resources that the attribute is associated to.

Defining Resource General Properties

To define resource general properties:

1. From the Model Workspace, right-click on a resource.

2. Select the General tab.

3. Complete these fields:

Note. The fields that appear on the page are dependent on the type of resource.

Code	Specify a name or a unique code for the resource.
Notes	Specify any supplemental information about the resource. This field is optional.
Last Operation Run	Specify the last operation run for the resource. This value is used to facilitate immediate changeovers, if required. This field requires the definition of an operation that uses this resource. This is an optional field and does not appear for suppliers or items.
Attached Document	Specify the path to document providing additional information. This field is optional and does not appear for suppliers or items.

Defining Resource Capacity

This section provides and overview of resource capacity and discusses how to define resource capacity.

Understanding Resource Capacity

Resource capacity is the potential amount of production that a resource can complete. Resource capacity is defined differently for each resource category.

- For crews and tools, capacity indicates the number of simultaneous operations that can be run on that resources.

- For machines, capacity indicates the number of simultaneous functions that can be performed by that resource when available.

  Machines can also be designated as Batchable.

  See "Batchable Resources," Creating Production Scheduling Models, Defining Resources, Understanding Resources, *Production Scheduling*

You can also indicate the dynamic capacity of each resource. This feature enables you to specify when and to what level the capacity of the resource might vary, due to varying crew sizes, seasonality, or regularly scheduled maintenance.

Defining Resource Capacity

To define resource capacity:

1. Access the Resource Properties window.

2. Select the Capacity tab.

3. Complete these fields:

Single	Select this option for single capacity resources.
Multiple	Select this option for multi-capacity resources and specify the Maximum field.
Batchable	Select to indicate that this resource is a batchable resource. When this radio button is selected the Minimum, Maximum, and Pull Forward Window fields are active. This option is only available on machines, not crews and not tools.
	Note. If this option is selected and this resource is a member of a resource set, all members of the resource set must also be Batchable.
	See "Batchable Resources", Creating Production Scheduling Models, Defining Resources, Understanding Resources, *Production Scheduling*
Minimum	Enter the minumum spot capacity limit. The system enters a value of 1 by default. The minimum can not be greater than the maximum. This field is only available when Batchable is selected as the Capacity Type.
	The minimum spot capacity is a soft constraint, which means that if there is not enough demand within the Pull Forward Window to fill the resource to the minimum, it may be violated. Supporting alerts are provided to inform the user when this occurs.
Maximum	Enter the maximum spots for this resource if the resource is Batchable. The system enters a value of 2 by default.
	Enter the maximum number of resources in this pool if the resource type is Multi-Capacity.
Pull Forward Window	Enter the period of time that the solver is permitted to pull operations that have the same attributes forward to fill the resource to at least the minimum batch capacity. The system enters a value of 7 days by default. This field is only available when Batchable is selected as the Capacity Type.
Dynamic Capacity	Select this option for multi-capacity resources with a varying capacity level. Complete the Date and Value fields.
Date	Enter the date when the capacity changes.
Value	Enter the revised maximum capacity. The value must be less than or equal to the Maximum Capacity value. Click Update to revise the capacity or Delete to delete a capacity entry.

Defining Throughput

This section provides an overview of throughput and an example of throughput rate, and discusses how to define throughput for a resource.

Understanding Throughput

The purpose of defining throughput is to adjust operation time to complete the operation in a given interval. For example, if the speed of a line changes on a daily or weekly basis based on available labor, operation time can be increased or decreased.

Production Scheduling enables a user to define dynamic throughput percentages on a given resource for a specific effective date. It also recognizes when throughput is faster or slower when creating a schedule.

On a resource, the throughput is defined as any percentage greater than or equal to 0 (zero) percent. When this resource is utilized in an operation (whether it be in a demand driven or work order driven model), the speed at which the operation runs is dynamically modified based on the defined throughput percentage.

A resource may also run at a rate which is greater than 100 percent throughput.

Example of Throughput Rate

This example assume that it takes 15 minutes to make a batch of 25 units of Item A, which achieves a 100 percent throughput rate.

This diagram illustrates this manufacturing process:
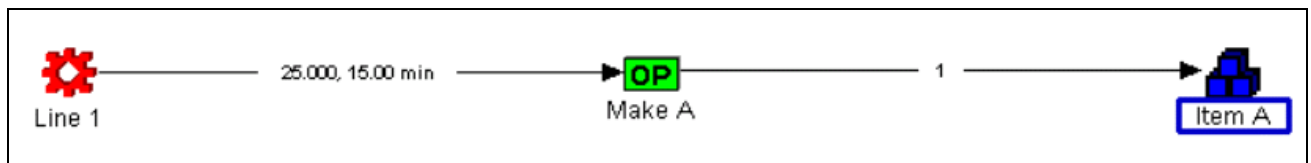
Throughput rate 100 percent

Assume that the user determines that in week two of the schedule horizon, line 1 will run at a 75 percent throughput rate. Also, an additional order for 2000 units is required on the Friday and the operation will be run JIT.

The first thing that the user must do is access the resource properties for Line 1 and select the Throughput tab. The user selects the Change Throughput Percentage icon. Upon doing this a blank row is added and the user adds the effective date, a description, and the 75 percent throughput value for the resource.

Resource Properties - Throughput page

The user then navigates to a schedule view, performs a solve, and begins analyzing the schedule using the combined Resource Gantt, Operation Gantt, Item Graph. Looking at the item graph, the rate at which Item A is being produced at the end of week 2 is at a slower rate. It is being produced at a 75 percent throughput rate.
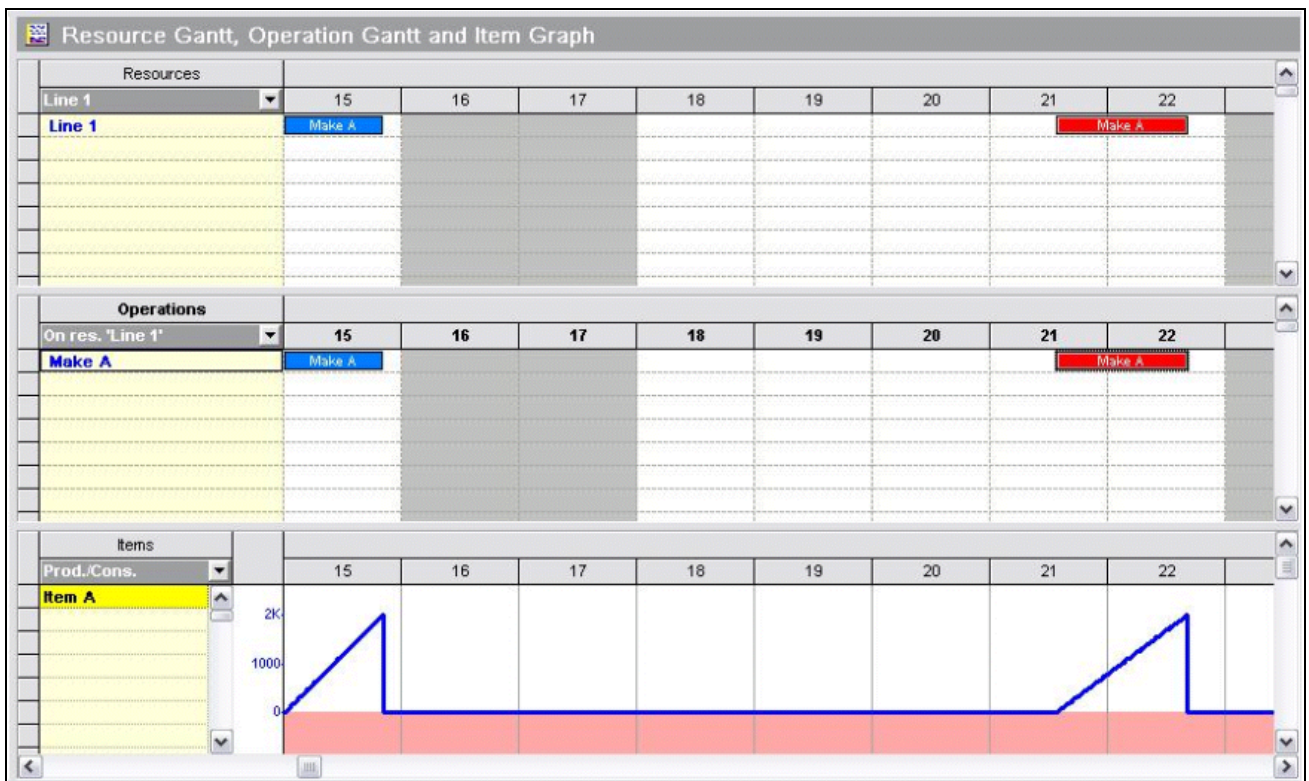
Throughput rate 75 percent

Defining Throughput

To define throughput for a resource:

1. Select the resource.

2. Right-click and select the Properties option.

3. Select the Throughput tab.

4. Click the Change Throughput Percentage button.

   A blank row is added to the page.

5. Complete these fields:

   Effective Dates   Enter the effective date of the schedule horizon.

   Description   Enter a description that explains the change in throughput.

   Throughput %   Enter the percent at which you want this line to run. For example, to run at 75
                                percent, enter 75.00%.

6. Click OK.

Defining Resource Solver Options

Solver Options apply to crews, machines, and tools.

To define resource solver options:

1. Access the Resource Properties window.

2. Select the Solver Options tab.

3. Complete these fields:

Minimize Changeovers	Select this option if changeover minimization is important for this machine, and if this machine is not considered a campaign run optimized resource. Otherwise, select Minimize Work In Progress.
Use Prebuild Maximum	This option is available if you have selected Minimize Changeovers. When meeting a demand, prebuild maximum is the amount of time prior to the point in time that the resource is required to fulfill the demand that the solver will look to minimize changeovers on this resource. For example, if you specify seven days, the solver considers the operations that are running that resource for the past seven days when minimizing changeovers.
Minimize Work In Progress	Select this option if changeover minimization or campaign run optimization is not required for this resource. The solver considers changeovers if specified, but the changeovers are not considered by the changeover minimization algorithm.
Use Campaign Run Optimization	Select this option to configure this resource to operate in defined run lengths. This option is only available for single-capacity resources, and the Enable Campaign Run Optimization option must be selected in the Model Properties window. Campaign Run Optimization requires a separate license.
	See "Setting Up Campaign Run Optimization," Configuring Solver Options, *Production Scheduling*
Minimum Cycle Time	Enter a value and select the time measurement required to cycle through demand for the items that are produced. Time measurement options are: *Shift, Days,* and *Weeks.*
Minimize Changeovers between cycles	Select to indicate that you want Production Scheduling to consider the operations that ran in the previous cycle. If possible, the application schedules an operation that can run on this resource without a changeover.
Offload Threshold %	Enter the percentage of capacity above which Production Scheduling offloads production to alternate resources in the set.
Use Prebuild Target	Select this option to check utilization when making off-loading decisions
Prebuild Target	Specifies the amount of time prior to the point in time that the resource is required to fulfill the demand on time. The prebuild target is used to check utilization when making off-loading decisions. You must specify a prebuild target value if this option is selected. Values are: *Minutes,Hours,Shifts,Days, Weeks.*
	Resource off-loading enables you to dictate resource load levels prior to the utilization of an alternate resource within a set. Using the Prebuild Target option, you can define the timefence.
	This option is not enabled automatically. It must be explicitly set. To determine the best value for your needs, you need to experiment with different values and understand the impact on the quality of your schedule.

Use Upstream Operation Buffering	Select to indicate that the solver should attempt to use upstream operation buffering for the resource.
	In some manufacturing environments, it may be desirable that a resource continues to run for as long as possible. This behavior may be dictated by a physical property of the device, such as a clean-out or purging that must occur after the resource has stopped running for some period of time. Select this option to indicate that the solver may place operations in continuous order so that the resource runs for as long a possible.
	Production Scheduling strives to create a highly optimized schedule, meeting demands as closely as possible to their request dates in order to avoid unnecessary inventory buildup and carry costs.
Target Buffer Size	Enter a value to indicate that the desired minimum amount of operations that the chosen resource should strive to run continuously before starting. This option is used in conjunction with the Use Upstream Operation Buffering option.
Relax Calendar and Capacity Constraints	Select this option to relax all constraints on this resource including any capacity and calendar associated with this resource.
Capacity Constrained Resource	Select this option to designate this resource as a capacity constrained resource. You must select the Use Theory of Constraints Principles option in the Model or in a schedule before you use this option.
	See "Setting Up the Scenario Properties - General page," Configuring Scenario Properties, *Production Scheduling*
Minimize Resource Idle Time when using Prebuild	Select this option to minimize idle time.
	When using Prebuild as a scheduling strategy, sometimes the sequence of operations to fulfill demands results in idle time on a resource. Conversely, the best sequence of operations to minimize idle time may not necessarily fulfill demands on time.

Defining Resource Availability

Availability applies to crews, machines, and tools.

To define resource availability:

1. Access the Resource Properties window.

2. Select the Availability tab.

3. Complete these fields:

Calendar	Select an applicable calendar. If no suitable calendar entries appear, click Open Calendar, and create an appropriate calendar for this resource.
Fixed Timefence	For suppliers, enter the fixed amount of time beginning at the start of the scheduling horizon that precedes any supplier lead time. For all other resources, enter the time period that the schedule is to remain fixed for this resource.

A timefence is a period of time when operations are not altered for a crew, machine, or tool resource. During the initial solve or subsequent repairs, the solver is only given the freedom to schedule other operations after the timefence ends. Any work order operations that are active or have a firm status of Firm Start, Firm End, or Firm Range and fall within this timefence may be scheduled in within the fixedtime fence.

Timefences are established to maintain the current schedule for critical resources so that they are not affected by any scheduling changes before the next schedule is finalized. A timefence is depicted in a schedule view using a red line in the Gantt view.

Note. When the schedule is modified and solved again, the operations that are assigned to this resource remain the same as the last schedule during the fixed timefence period.

Enforce Horizon Start	Select to move forward all operations running behind on various machines, without changing the sequence of the operations and consider upstream or downstream dependencies.

There are a variety of different scenarios associated with this feature in regards to why operations may appear running in the past. For example, work order operations that have a Firm Start Date, Firm End Date, or Firm Date Range in the past, several work order operations that have an active status on the work order routing operation step that conflicts with another work order, and so on.

Defining Resource Costs

Resource costs are costs incurred when a resource is used and can be defined for a variety of resources including crews, items, machines, and tools. Machine costs are used by the Campaign Run Optimization (CRO) algorithm, which is not the standard base solver in Production Scheduling. Machine, crew, and tool costs are used for reporting purposes if the CRO algorithm is not being used. Carrying and stock-out costs are recorded for items. Setup and operating costs are gathered for all of the other resources. The base Production Scheduling solver does not use these costs.

Note. The period that is specified on the cost tab is set in system options. From the Tools menu, select Options and click the Global Settings tab. From the Cost Time Unit field, the choices are minute, hour, shift, day, and week.

To define a machine, crew, or tool resource costs:

1. Access the Resource Properties window.

2. Select the Costs tab.

3. Complete these fields:

Setup	Specify a fixed cost that is incurred each time that the resource is used. This cost applies to crews, machines, and tools only.
Operating	Specify a cost that is incurred over time for the duration that the resource is used. This cost applies to crews, machines, and tools only.

Defining Resource Attributes

Use this page to inform the Production Scheduling solver which attributes should be used on this resource when grouping operations of the same type together. This is essentially creating windows of attribute types.

To define machine resource attributes:

1. From the Model Workspace, right-click on a machine resource.

2. Select Properties.

3. Select the Attributes tab.

    Available attributes appear on the left side of the page.

    Note. The Attributes tab is only available if the Batchable radio button is selected on the Capacity tab.

    See "Defining Resource Capacity," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

4. Select an available attribute and click the right arrows to move an attribute to the Assigned Attributes column.

5. Set the changeovers as Concurrent or Sequential if there are multiple attributes

Concurrently	In this case, the maximum time and maximum cost is the applicable changeover rule.
Sequentially	In this case, the sum of both time and cost is relevant from a changeover perspective.

    It is important to note that in many cases, the traditional manner in which changeovers are evaluated is via the use of wildcard "else" rules and is represented usually be a "*" to and from rule on the machine. In this case, the rules would likely not want to be summed in terms of time or cost, so these should not be included in the concurrent or sequential calculation IF other rules are applicable.

6. If multiple calendars are applicable for a changeover, select the availability by selecting an option below:

Any calendar may be up	This option is interpreted as the union of the calendar. Any Calendar can be "up" for the changeover to occur.
All calendars must be up	This option is interpreted as the intersection of the calendar. All Calendars must be "up" for the changeover to occur.

The calendar uptime is illustrated below.
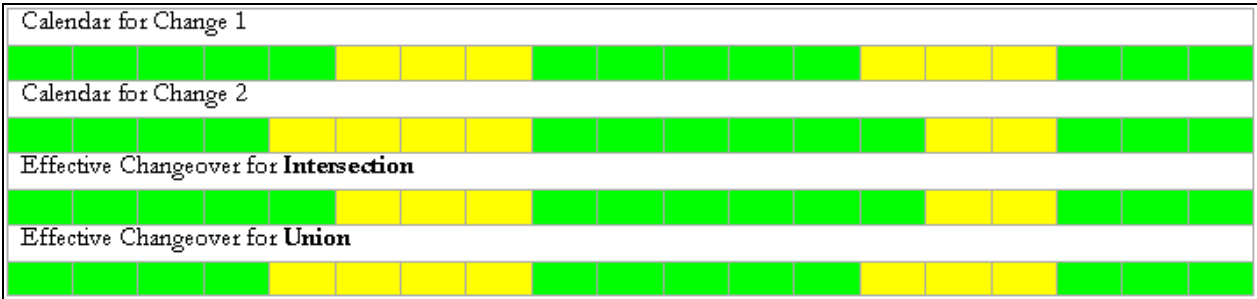
Uptime calculation

7. Click OK to save your changes.

Changing Multiple Reusable Resources

In Production Scheduling, you can change multiple reusable resources at one time to change solver options, availability, and costs settings.

To change the properties of multiple resources:

1. Access the Model Workspace.

2. In the Model Workspace, select the resources that you want to change by holding down the Shift key and highlighting the appropriate resources or folders. If the resources or folders are not listed sequentially, hold down the CTRL key and select the individual resources.

    Note. If you select a folder, all resources in that folder are considered.

3. Right-click and choose Change from the menu.

4. Select the Solver Options tab and adjust the necessary solver settings.

5. Select the Availability tab and adjust the availability options.

6. Select the Cost tab and adjust the setup and operation cost settings.

7. Click OK to save the changes and close the window.

8. To view the results of the changes, click on one of the modified resources.

Defining Item Constraints

You can track item and shared storage space constraints including minimum and maximum quantities of items, and shared storage space resources.

The Constraints tab is available for only item and shared storage space resources. When solving, the system attempts to create a schedule that meets the item levels as closely as possible. The minimum stock value is a hard constraint for work in process components and a hard constraint for items with order requirements. Safety stock is a soft constraint that can be violated if necessary. Safety stock should be replenished immediately after a safety stock violation occurs.

Use this page to define item constraints:

Item Properties - Constraints page

To define item constraints:

1. From the Model Workspace, right-click on an item and select Properties.

2. Select the Constraints tab.

3. Complete these fields:

Unit of Measure	Select a unit of measure for the item. All units of measure are available from the global list.
Display Precision	Specify the number of decimal places. For example, the number 5.125 has three decimal places so the precision is three. Setting a precision value is required when items are not measured using whole numbers (for example, by weight).
Minimum	For items only, specify the minimum amount of this item to be carried in inventory. If the minimum amount of this item is going to be violated while solving the schedule, the system tries to delay the operation as long as possible to allow replenishment. If restocking does not happen, then the solve fails, generating an error message.
Maximum	For items only, specify the maximum amount of this part to be carried in inventory. If the item is exceeded, an error is generated, and the solve is discontinued.
Relax Minimum	Select this option to ignore minimum constraint values for this resource during a solve. Relaxing this constraint allows the values for this item to fall below the defined minimum. Replenishment activities are not planned.
Relax Maximum	Select this option to ignore minimum constraint values for this resource during a solve.

If you relax the minimum and maximum constraint, you should monitor the item graph for potential violations. |
| **Start** | Specify the amount of inventory at the start of production.

This value represents the beginning inventory for this item at the factory for which a schedule is created and may consist of multiple lots. |
| **Storage Space** | Select a storage location. An item can be assigned to only one shared storage space. The assigned storage space must have the capacity to store the initial inventories for all items allocated to it; otherwise an error message is generated.

This is an optional field and should only be populated if storage represents a critical scheduling constraint. |
| **units of item per** | Specify a value that represents how much of one storage unit of space (for example, a pallet, cubic meters, or other unit of measure) is used for one unit of item. For example, if a tote can store 100 units of this item, then *units of item per* the storage is represented as .01. |
| **Date** | Specify the starting date for the safety level using the small arrows or the calendar. This field is optional, and applicable to when you define safety stock levels. If you want to always have a safety level, set this date to the beginning of your schedule horizon. |

Value	Specify the item quantity (a number less than the number in the Maximum field) that represents the desired safety stock. You can input varying safety levels (for example, to account for seasonal variations) if required. Click the Update button. The system displays the entry in the Date Value field. You can enter multiple values. The system uses the first safety stock entry until the next entered date.

Defining Item Costs

Item costs are only considered by the Campaign Run Optimization (CRO) solver engine and are also used for reporting Key Performance Indicators. The base solver does not consider these costs when creating a production schedule. If you are using the CRO algorithm, these costs must be defined.

Use this page to define item costs:

Item Properties - Costs page

To define item costs:

1. From the Model Workspace, right-click on an item and select Properties.

2. Select the Costs tab.

3. Complete these fields:

Carrying	Specify the unit cost of carrying this item for the period. This cost applies to items only.
Stock-out	Specify the unit cost that is associated with running out of stock for this item during the period. This cost applies to items only.
Safety Stock Violation	Specify the cost that is associated with violating the safety stock. This cost applies to items only.

Defining Item Attributes

To define item attributes:

1. From the Model Workspace, right-click on an item and select Properties.

2. Select the Attributes tab.

3. Complete these fields:

Manufactured	Select this option if the item is manufactured.
Purchased	Select this option if the item is purchased.
Saleable	Select this option to indicate that the item should be available on the item pick list in the Supply and Demand editor.

4. To assign an attribute-base resource to an item, select the check box to the left of the Name field and, if necessary, assign the appropriate attribute value from the Value list.

    For example, select Duration or Temperature. Any other items sharing the same attribute assignments are eligible to share the same batchable resource.

    The Attribute Type and Attribute Value fields are used in conjunction with batchable resources and changeover rules. To populate these fields, you must define Attributes in the Model Workspace under the All Resources heading. For example, an attribute type might be Temperature and have several values such as 500 Degrees, 450 Degrees, and so on. Once these are defined, you can associate them to items. Any items that share the same attribute assignments are eligible to share the same batchable resource and changeover rules.

    See "Batchable Resources," Creating Production Scheduling Models, Defining Resources, Understanding Resources, *Production Scheduling*

Defining Item Suppliers

All items in the model must have a source of supply that can be either a manufacturing operation or a supplier item lead time. For purchased items, it is recommended to setup a supplier representing a lead time in which the item can be procured, and a large order multiple. This creates a material constrained schedule.

To define item suppliers:

1. From the Model Workspace, right-click on an item and select Properties.

2. Select the Suppliers tab.

3. Complete these fields:

Receiving Calendar	Select a receiving calendar for the item. By associating the item with a calendar, the item can not be received during certain calendar events, for example during a weekend. If no suitable calendars appear, click the Open Calendar editor, and create an appropriate receiving calendar for the item.
Supplier	Select a supplier to associate with this item if it is a purchased item. Suppliers must be defined as a resource in the model before they appear in this list.
Preference	Assign a supplier preference to this supplier. This value must be a positive integer of 1 or greater. This value is used in the event that multiple suppliers can supply the same item.
	Production Scheduling selects the preferred supplier unless the secondary supplier has a shorter lead time, and the material is needed within this shorter lead time.
Lead Time	Specify the lead time necessary to supply the purchased item.
Order Multiple	Specify the standard order multiple for the purchased item. For example, items must be purchased in batches of 1,000 units.
Buffer	Select a buffer value for the supplier if you are using Theory of Constraints principals to create your schedule. This value represents the time between the supply event and the start of the operation that uses the item. This value is used when Production Scheduling calculates purchase order arrive dates. You must create a buffer before you can select it from this drop-down list box.

Changing Multiple Items

In Production Scheduling, you can change multiple items at one time to change constraints, costs, attributes, and suppliers.

To change the properties of multiple items:

1. Access the Model Workspace.

2. In the Model Workspace, select the items that you want to change by holding down the Shift key and highlighting the appropriate items or folders. If the items or folders are not listed sequentially, hold down the CTRL key and select the individual items.

    Note. If you select a folder, all items in that folder are considered.

3. Right-click and choose Change from the menu.

4. Select the Constraints tab and adjust the necessary minimum and maximum levels.

5. Select the Costs tab and adjust the carrying, stockout, and safety stock violation settings.

6. Select the Attribute tab and adjust the attribute-base resources assigned the items. Select the check box to the left of the Name field and, if necessary, assign the appropriate attribute value from the Value list.

7. Select the Suppliers tab and adjust the receiving calendar and the supplier attributes.

8. Click OK to save the changes and close the window.

9. To view the results of the changes, click on one of the modified items.

Defining Resource Groups

Resource groups enable multiple resources to be combined as a group when viewed in schedule views or represented in the Changeover editor (for machines or tools only). Organizing resources into logical groupings can help you to organize your model data. For example, you might want to organize resources by resource type or vendor. Similar to resources, resource groups can be organized in folders. Additionally, resources can reside in more than one resource group.

You can set a default resource group when you first open schedule views. When a schedule view is opened, only those resources in the default group appear in the schedule. Each schedule view can have its own default group, with the currently selected default appearing with a dot beside the name.

To define resource groups:

1. Access the Resource Group Properties window.

2. Specify a new name for the group in the Code field.

3. Select the Attributes tab.

4. In the Used in Views section, select the schedule views for which this resource group will be available.

    For groups that consist of only machines or tools, you can also select the Changeover editor.

5. Click Apply, and then click OK.

6. Click the plus sign (+) beside All Resources and each folder to display all resources.

7. Drag and drop one or more resources into the new group.

    You can also add the same resource to different groups. Press Shift or Ctrl to make multiple selections.

Defining Visits

This section provides an overview of visits and discusses how to:

•   Add a visit.

•   View and define visit properties.

•   Create milestones.

•   Modify milestones.

•   Delete milestones.

Understanding Visits

The scheduling of maintenance resources at a maintenance depot can be quite complex and presents a significant challenge to customers looking to efficiently schedule these resources. Production Scheduling provides the following constructs to provide this functionality:

- Visits – This new data construct allows users to associate to one or more work orders, where the work orders will inherit a Visit "earliest state date" such as the anticipated aircraft arrival into a maintenance depot and the request date associated with the completion of this visit.

- Milestones – Provide the ability to specify logical grouping of work orders within a Visit. The combination of Visits and Milestones provide an advanced capability to define task hierarchies and logically and flexibly group work orders together in both concurrent and sequential patterns where work can diverge and converge.

- Work Breakdown Structure – Several new key attributes are added to work orders which can be used for organizing and filtering work orders.

Scheduling visits must take into consideration numerous constraints such as complex task hierarchy relationships, logical groupings of work, maintenance capacity, calendars and more. This becomes increasingly difficult when multiple visits are considered simultaneously and have competing priorities. Customers with maintenance scheduling problems have expressed the requirement to be able to be able to schedule visits taking into account:

- Visit start/end dates

- Visit milestones

- Task hierarchies

- Crew, machine, tool, and material constraints

- Operation/work order precedence constraints

- Grouping of work based on work breakdown structure

While the maintenance scheduling problem may be applicable to many industries, the focus of the solution is solving short term aircraft maintenance scheduling problems.

See "Scheduling Maintenance Work Orders" for more information.

Adding a Visit

To add a visit:

1. Right-click All Visits from the Model Workspace and select New Visit. A new visit appears under All Visits. A visit appears under the folder with a system-generated name.

2. Enter the appropriate name for the visit and press Enter.

    Visits are sorted alphabetically under the All Visits folder.

Viewing and Defining Visit Properties

To view and define visit properties:

1. Expand the All Visits folder.

2. Right-click the visit you want to view or define and select Properties.

3. From the visit Properties – General tab, review or modify the following fields:

Code	Enter the visit code. This is the name of the visit displayed in within Production Scheduling.
Earliest Start Date	Represents the earliest date the visit could be scheduled which would typically be the arrival of the aircraft. This is treated as a hard constraint.
Request Date	The request date of the visit.
Start Date	The projected start date of the visit as determined by the Production Scheduling solver based on the various constraints.
Completion Date	The Completion date of the visit determined by the Production Scheduling solver.
Notes	Enter relevant notes.

4. Select the Milestones tab to view the milestones assigned to the selected visit. To view milestone properties, select a milestone and click Properties. The Actual Date is the projected completion date of the Milestone determined by Production Scheduling.

5. Select the Work Order tab to view work orders that associated with this visit. Double-click a work order to view it in the Work Order Editor.

6. Click OK to apply and close changes made in the visit Properties window. Click Cancel to close the visit Properties window without saving your changes.

Creating Milestones

To create milestones:

1. Expand the All Visits folder.

2. Expand the visit folder. All milestones associated with the visit appear below the visit name in the Model Workspace.

3. Right-click on the visit and select Add Milestone. A milestone appears under the folder with a system-generated number.

4. Enter the appropriate name for the milestone and press Enter.

   Milestones are sorted alphabetically under the visit.

Modifying Milestones

To modify milestones:

1. Expand the All Visits folder.

2. Expand the visit folder. All milestones associated with the visit appear below the visit name in the Model Workspace.

3. To rename a milestone, right-click the milestone and select Rename. Enter a new name and press Enter.

4. Right-click on a milestone and select Properties.

5. From the milestone Properties – General tab, review or modify the following fields:

Code	Enter the visit code. This is the name of the visit displayed in within Production Scheduling.
Actual Date	The projected completion date of the milestone based on the Production Scheduling solver.
Notes	Enter relevant notes.

6. Select the Work Order tab to view work orders that are dependent or assigned to this milestone. Double-click a work order to view it in the Work Order Editor.

7. Click OK to apply your changes and close the milestone Properties window. Click Cancel to close window without saving your changes.

Deleting Milestones

To delete milestones:

1. Expand the All Visits folder.

2. Expand the visit folder. All milestones associated with the visit appear below the visit name in the Model Workspace.

3. Right-click on a milestone and select Delete. The milestone is removed from the Model Workspace.

Defining Operations

This section provides an overview of operations and discusses how to:

• Create an operation diagram.

• Designate a primary output.

• Define operation type.

- Define operation costs

- Define effective dates

- Define item requirements

- Define requirements for crews, tools, and machines.

- Create item sets.

- Create resource sets.

- Create all of sets.

- Change multiple operations.

Understanding Operations

This section discusses:

- Operations

- Operation groups

- Duration

- Lot sizes

- Primary outputs

- Identify operations

Operations

An operation is a production event that consumes and produces items, and uses crews, machines, tools, and vendor capacity. Operations establish associations between different resources. Operations combine to form routings that are used to produce specific items. An operation can have various inputs and outputs, depending on the complexity of your production process. You must model all of the operations in your production process in order to generate a production schedule using Production Scheduling.

Operations are defined by the resources that they require. Operations can transform materials: they can consume one or more items when they start, and produce one or more items when they complete. Operations use resources and have durations. All operations have the following characteristics:

Resource requirements	An operation can have capacity requirements for a resource that it needs. An operation can require more than one resource, or one or more of a subset of resources. It can also consume alternate resource sets, for example: crew 1 and machine 1, or crew 2 and machine 2.
Item consumption	All items are consumed when the operation starts. The consumption is a fixed amount that is based on the duration of the operation. An operation can have the flexibility to select from a set of items.

Item production   All items are produced when the operation completes. Production can be a fixed amount or dependent on the duration of the operation. Production rates can also depend upon the resources that are involved.

Costs   Costs can be fixed setup costs or production costs. These costs are used for reporting purposes only.

Once you have set up resources and operations, you can create a diagram by dragging resources from the Model Workspace into an Object Window. If you create item or resource sets, you must also drag buttons from the Resource toolbar into the diagram. All resources are available for use except the shared storage spaces and Attributes.

An operation diagram is a graphic representation of your production environment that enables you to see relationships and make adjustments where necessary. The model also enables you to change the properties of objects so that the application can take into account every detail in your model.

An operation diagram is used to detail all of the resources that are required to complete the operation and delineate which resource specifies the duration. The capacity and amount of resources that are consumed or produced are also modeled. The operations are created and organized in the Model Workspace. Logical folders are useful for maintaining model structure.

Create operation diagrams, or visual representations of your production environment, by dragging and dropping resources into the diagram, and then connecting them to the operation button. As resources are added to an operation, they appear beneath the operation's button in the Model Workspace.

Operation Groups

Operation groups can be created to assist with schedule analysis. Operation groups enable two or more operations to be combined as one within the Operation Gantt schedule view as well as the Changeover editor. For large models, operation groups filter the schedule views so that you can quickly review those operations, thereby saving time. Applying the same time for changeover to all group members reduces the number of rules required. Create as many groups as you need for your requirements. Operation groups can be organized in folders, similar to operations, if desired. Operations can reside in more than one operation group.

Duration

Every operation must have a defined duration, determined by only one resource or resource set. For example, if an operation uses a machine with a duration of one hour and a crew, Production Scheduling allots one hour for the crew. Using item and resource sets, different durations can be recorded for those resources in the set to complete the same operation. For example, one machine can take much less time to perform an operation than another machine.

This diagram illustrates an operation with a resource-dependent duration specified within a set:
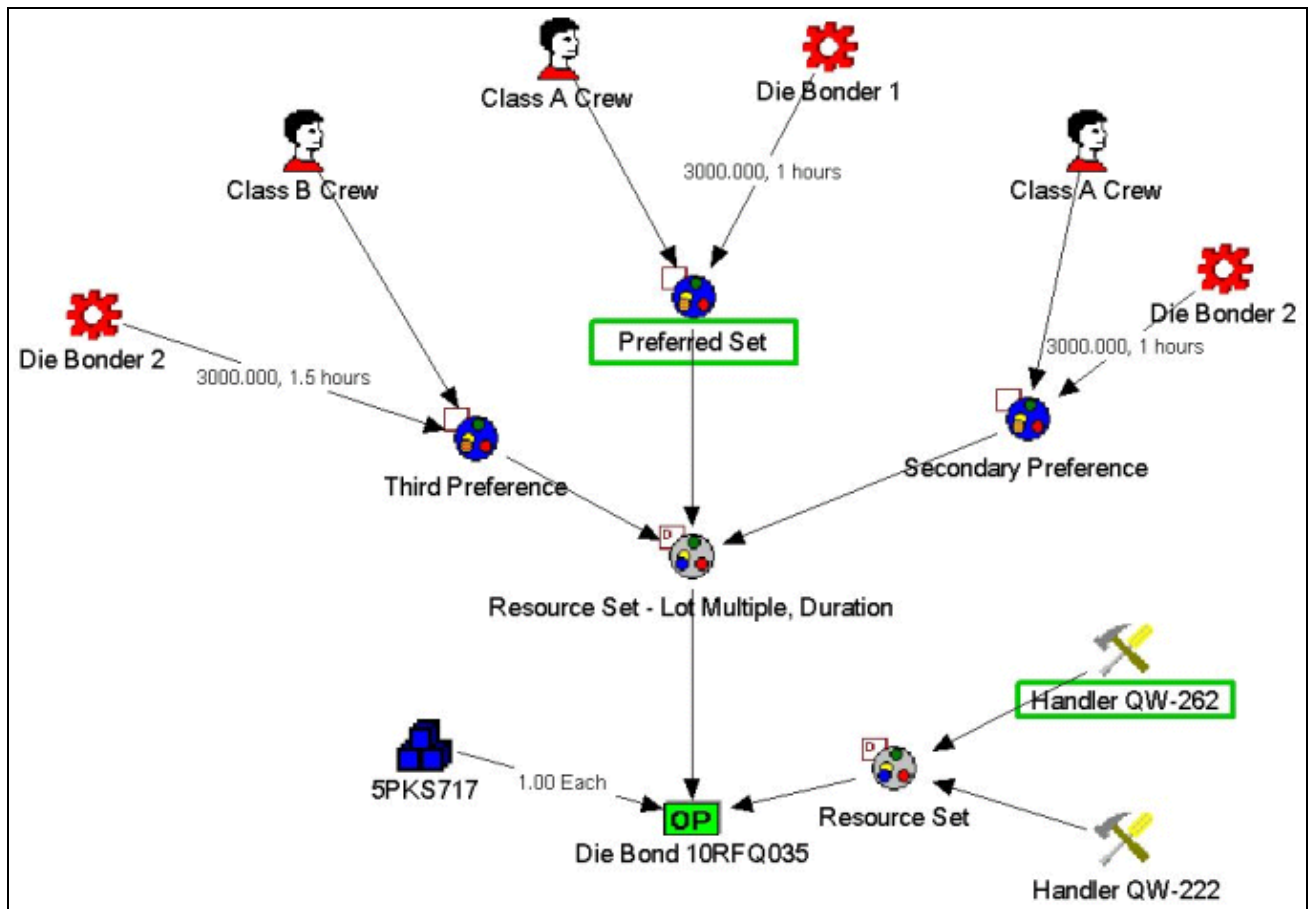
Resource Dependent duration within a set

The preferred set, which comprises of Class A Crew and Die Bonder 1 has a duration of one hour. The secondary preference, which comprises of Class A Crew and Die Bonder 2 takes the same amount of time. However, the third preference, which comprises of Class B Crew and Die Bonder 2 has a duration of 1.5 hours to make the same quantity. When run, either Handler QW-222 or Handler QW-262 can be used.

For duration to be attached to elements in a set, the set properties must first have duration selected. Then all of the machines within the set can be defined with a duration attribute. No other durations are permitted in this operation.

Lot Sizes

Whenever an operation produces more than one identical item, the resulting output is called a lot. For example, consider a stamping machine that produces eight identical items in a single operation. In this example, the lot size is eight.

You can use lots to reduce the total number of operations in your schedule. You can do the following:

• Increase the operation lot-size and duration.

• Define logical lot sizes if physical ones do not exist.

You create a logical lot size by increasing the quantity that an operation produces, as well as increasing its duration. For example, an operation that runs for one minute and produces one item can also be represented as a single operation that runs for ten minutes and produces ten items. By using logical lot sizes, you can dramatically reduce the total number of operations that the solver must schedule.

Note. To avoid excess inventory, lot sizes should be divisible into demand.

Although using logical lot sizes can help you reduce the number of operations in a schedule, it can also artificially increase the amount of time required to produce a specified number of items. Item and precedence relationships are still enforced and because each operation takes longer to complete, subsequent operations are also delayed.

This diagram illustrates the effects of using a logical lot size:
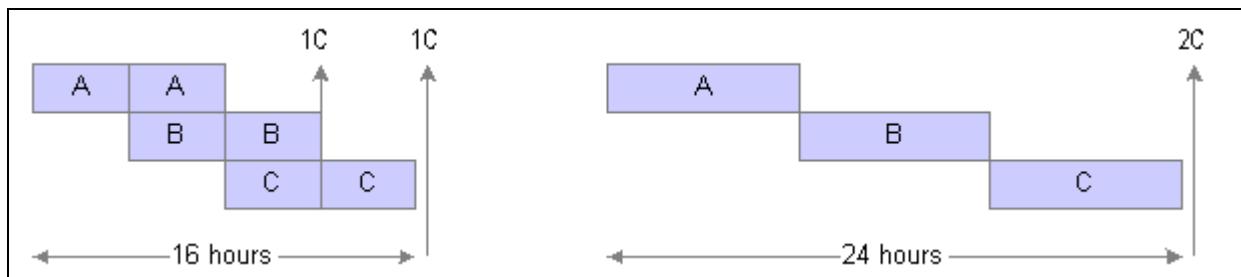
Logical lost size

In the first illustration, each operation has a duration of four hours and produces one item. Operations B and C have a precedence relationship of Start after Start. In this example, item C is produced at 12 hours and again at 16 hours. In the second illustration, the number of operations has been reduced using a logical lot size of two. Because operations B and C cannot begin until the previous operation has started, the makespan has been exaggerated to 24 hours.

Crew and machine resources in an operation can, optionally, have a lot multiple attribute. This multiple is applied to all items that are consumed or produced by that operation. A lot multiple is used to specify that an operation can consume or produce different amounts of materials if the operation runs on a different resource. For example, a certain press can have a larger template that enables it to stamp out twice as many items at a time as a smaller template. This situation has a lot multiple of 2. The lot multiple is applied in the same way to all inputs and outputs of the operation, but the lot multiple does not affect the duration of the operation.

Only one resource determines the lot multiple for an operation, although with the use of sets, more than one lot multiple might be specified. Most operations do not need to specify a lot multiple. If a lot multiple is not specified, then the lot multiple is assumed to be 1. With the use of sets, different resource lot multiples can be specified. Within the set configured for a lot multiple, each resource must have a lot multiple recorded. Production Scheduling selects a resource in the set and uses the lot multiple specified for that resource.

Operations can have both a duration and a lot multiple. The duration and lot multiple can depend on different resources. For example, the crew determines the duration of the operation, and the machine determines the lot multiple. Alternatively, they can depend on the same resource. For example, the machine can determine both the duration and the lot multiple of the operation.

In the event that an operation with an associated work order is offloaded to an alternate machine in a resource set, the item quantity produced by the primary machine will affect the duration of the operation on the alternate machine. The quantity of produced and consumed items will remain the same. The application will calculate the new duration on the alternate machine by applying the production rate of the alternate machine to the item quantity produced by the primary machine.

For example, machine A and machine B are a part of a resource set, and you want to offload an operation with an associated work order from machine A to machine B. The duration of the operation running on machine A is 1 hour, with a lot multiple of 100 units (100 units per hour). The duration of the operation running on machine B is 2 hours, with a lot multiple of 50 units (25 units per hour). The operation produces 1000 units and runs on machine A for 10 hours.

To produce 1000 units on the alternate machine that produces 25 units per hour, the new duration required is 40 hours.

Primary Outputs

To take advantage of the pegging capabilities of Production Scheduling, the primary output of the operation needs to be identified. Pegging enables you to view all of the item resources and operations associated with each line item in a demand order, and how they have been allocated. In addition, the Resource, Demand, and Operation Gantt views can display all of the resources and operations related to a specific operation "upstream" (backward) or "downstream" (forward).

The primary output of an operation can be a specific item . If an operation does not produce any items, a primary output does not need to be designated.

Identifying Operations

All operations must be identified within Production Scheduling in detail before feasible schedules can be generated. As with resources, when operations are created manually within the application, they are started in the Model Workspace and can be organized with the use of folders.

As with resource folders, operation folders are designed to organize your operations. The organization of operations is defined by you. Production Scheduling does not draw any implications or conclusions from the order. The All Operations and All Resources folders appear in the Model Workspace by default.

Before adding an operation, ensure that all resources have been entered. As with resources, establish a naming convention that clearly and uniquely identifies each operation. Distinctly named operations are helpful when analyzing schedules. Operations in the XML document also have an identifier called *integrationId,* which can be used to ensure uniqueness. The code for the operations can then be named something that is meaningful to the planner.

Creating an Operation Diagram

To create an operation diagram:

1. Access the Model Workspace.

2. Double-click an operation.

3. Select the operation button with the operation name appears in a window.

4. Below your model, click the plus sign (+) beside All Resources to expand all of the relevant folders.

5. Click a resource to be added, and drag it into the Object Window.

   Repeat for each resource.

6. From the Resource toolbar, click the Insert Link button.

7. In the operation diagram, click the source button once (produced resource), and then click the destination button (consumed resource).

    An arrow appears between the source and the destination with an empty square box in the middle of the line (the resource relationship). When a resource is linked to an operation, the property page for the created arc appears.

8. Specify the following link properties, and then click OK:

    Capacity Required   Specify the capacity required to run this operation, or accept the default value of 1.

    Lot Multiple   Select if the resource uses a lot multiple. A default lot multiple value of 1 is displayed. If you need a larger lot multiple, specify a different value. The value in this field is applicable only if the operation produces an item.

    Duration   Select if the resource is determining the duration. If you select this option, select a time unit for the duration. The time unit can be days, hours, minutes, or seconds.

    Minimum Run Length -   Select to specify the size of the production run in terms of quantity of items
    Quantity   produced. The value in this field is applicable only if you are using Campaign Run Optimization and the resource type is *Machine.*

    Minimum Run Length -   Specify to specify the size of the production run in terms of run time. The
    Run Time   value in this field is applicable only if you are using Campaign Run Optimization and the resource type is *Machine.*

9. Repeat the previous step to link all of the resources with the operation.

Note. To have the arrow point in the correct direction, always begin with the source button and end with the destination button. The operation is usually the destination button with one exception: if any items are being produced, then the operation is the source, and the new item button is the destination.

Designating a Primary Output

To designate a primary output:

1. Access the Operation Diagram.

2. Click the item that you want to designate as the primary output of the operation.

3. Select Primary Output.

Defining Operation Attributes

There are three activity types of operations: Production, Setup, or Cleanup.

Production operations are broken down into units of effort, however, setup and cleanup operations are not broken into units of effort when work orders are scheduled according to their units of effort.

See "Understanding Scheduling Work Orders According to Units of Effort," Configuring Solver Options, Setting Up General Solver Options, *Production Scheduling*

Example of Setup and Cleanup Operations

A routing produces Item A in a lot multiple of 25 units. The routing includes three operations:

*   Operation A, which takes 3 hours.

*   Operation B, which takes 2 hours.

*   Operation C, which takes 4 hours.

An order for 75 units of Item A is due on January 11th, at 4:30PM. It is scheduled as specified on the work order for a total of 27 hours:

*   Operation A for 9 Hours.

*   Operation B for 6 Hours.

*   Operation C for 12 Hours.

If using the work order units of effort feature, the work order is scheduled how it is executed on the shop floor:

*   Inventory is moved between operations in 25 unit increments.

*   Makespan is reduced by 10 hours, making the routing run time 17 hours.

See "Understanding Scheduling Work Orders According to Units of Effort," Configuring Solver Options, Setting Up General Solver Options, *Production Scheduling*

If operation A is a setup operation and operation C is a cleanup operation, they only need to be executed once as opposed to three times. Therefore, the setup and cleanup operations only run once per routing. The production operation runs three times, which is according to the routing lot multiple. This reduces the routing run time to a total of 13 hours:

*   Operation A runs for 3 hours.

*   Operation B runs for 6 hours.

*   Operation C runs for 4 hours.

Defining Operation Attributes

To define operation attribute:

1.   Access the operation Properties window.

2.   Select the Attributes tab.

3.   In the Type field, select one of these options: *Production Activity,Setup Activity,* or *Cleanup Activity.*

4. To define effective dates, complete these fields:

Effective Dates	Select this option to define effective dates for this operation.
Valid from	Specify the date when the operation becomes valid. The default value for this field is the beginning of the schedule horizon.
Valid to	Specify the date when the operation ceases to be valid. The default value for this field is the end of the schedule horizon.

5. To assign an attribute-base resource to an operation, select the check box to the left of the Name field and, if necessary, assign the appropriate attribute value from the Value list.

Defining Operation Costs

Operation costs are for reporting purposes only.

To define operation costs:

1. Access the Properties window.

2. Select the Costs tab.

3. Complete these fields:

Setup	Specify the one-time cost that is incurred every time that this operation is run (if applicable).
Operating	Specify the cost of running this operation for a specific period of time.

Defining Item Requirements

When defining item requirements, ensure that the material input quantity specified is according to 1 unit of output of the primary item of operation or the routing. The solver multiplies the quantity specified by the lot multiple, which is specified on the operation resource where lot multiple is specified.

To define item requirements:

1. Access the Link Properties window.

2. Complete this field:

Quantity	Specify the number of item units that are produced or consumed by one instance of this operation. The unit of measure that is displayed to the right of the quantity is inherited from the item's property information.

Defining Requirements for Crews, Tools, and Machines

To define requirements for crews, tools and machines:

1. Access the Link Properties window.

2. Complete these fields:

Capacity Required	Specify the capacity required to run this operation, or accept the default value of 1.
Lot Multiple	Select if the resource uses a lot multiple. The system displays a default lot multiple value of 1. If you need a larger lot multiple, specify a different value. The value in this field is only applicable if the operation produces an item.
Duration	Select if the resource is determining the duration. The duration should be specified in accordance to make the defined lot multiple. If you select this option, select a time unit for the duration. The time unit can be days, hours, minutes, or seconds.

Creating Item Sets

You can define and insert alternate items known as item sets. Item sets (a collection of similar resources) can be used to represent alternatives for items. Production Scheduling selects one component from each set when the operation is scheduled. You can create one or more item set.

Sets allow choices to be made that are based on the availability of, and preference for, the operation. Lot multiple specifies the quantity of items produced by an operation while duration specifies how long the operation runs. All operations must have a duration.

To create a item set:

1. Access the Operation Diagram window.

2. On the Resource toolbar, click Insert Item Set to create an item set.

3. Click an open area of the operation diagram.

   This process creates an empty set. The set button appears in the operation diagram.

4. On the Resource toolbar, click Insert Link.

5. In the operation diagram click the set button, and then click the operation button (in the source-to-destination order).

   The set is now linked to the operation, as indicated by a connector line with an arrow pointing to the set.

   Note. To return the cursor to a pointer at any time, right-click your mouse.

6. On the Model Workspace, drag and drop any item resources into the diagram.

7. On the Resources toolbar, click the Insert Link button.

8. In the operation diagram, click a resource, and then click its set button.

Note. Right-click anywhere in the operation diagram to return the cursor to a pointer.

---

9. To define the properties of resources that are attached to the set, do the following:

10. Right-click the link between the set and resource button, and select Properties.

11. In the Preference field, enter a value between 1 and 100 to indicate the priority that you want assigned to this resource.

    The solver selects the resource with the lowest priority, if possible.

12. In the Capacity Required field, enter the number of that resource that is required to perform the operation.

13. In the Lot multiple field, enter a value to scale item quantities for the operation.

14. In the Duration field, enter how long the operation runs.

Creating Resource Sets

You can define and insert groups of resources known as sets, which can be a resource sets. Resource sets (a collection of similar resources) can be used to represent alternatives for crews, machines, tools, and vendors. Production Scheduling selects one component from each set when the operation is scheduled. Or, if you are using All of sets, it may include several inclusive resources in the set. You can create one or more set. Resource sets enable Production Scheduling to select the most efficient resource for any given point in time, based on preference, capacity, lot multiple, and duration.

Sets allow choices to be made that are based on the availability of, and preference for, the operation. You can also specify lot multiple and duration for resource sets. Lot multiple specifies the quantity of items produced by an operation while duration specifies how long the operation runs. All operations must have a duration.

Sets provide operations with alternatives. For example, Operation A can be designated to use the set of machines M1, M2, and M3. Operation B can be designated to use a different set that also includes the same machines: M1, M2, and M3. By using sets, the schedule has the flexibility of using the most efficient combination of machines M1, M2, and M3 for the two operations. Use item and resource sets whenever options exist for completing part of an operation. By giving the solver alternatives, you have more potential ways of manipulating the schedule.

Lot multiples can appear only once in an operation in either a non-item resource or in a resource set. For the lot multiple attribute to appear in a set, the lot multiple property must be selected in the set's properties. Then all of the elements in the set can be defined with a lot multiple.

To create a resource set:

1. Select the Insert "One of" Set button from the toolbar.

2. Move this resource set into the diagram.

3. Right-click on the resource set and select Properties.

4. Select the Attributes tab.

5. Select the Lot Multiple and the Duration check boxes.

6. Click the OK button.

7. Click the Insert button on the toolbar.

8. Click the "1 of" Resource Set icon and then click the operation icon in the operation diagram.

    Click the set and operation icons in the source-to-destination order. The set is now linked to the operation, as indicated by a connector line with an arrow pointing to the operation.

    Note. Right-click your mouse to return the cursor to a pointer at any time.

Creating All of Sets

Resource sets enable you to substitute resources within an operation such as substituting machine 1 for machine 2, or crew 1 for crew 2. *All of Set* sets enable you to select groups of resources within an operation, such as running an operation on line 1 and operator 1, or running the operation on line 2 and operator 2. In this case, when using *All of* sets, line 1 can not be run by operator 2, and line 2 cannot be run by operator 1. *All of Set* is a set type that means all resources in the set are used to run the operation.

Note. You must define all possible combinations of resources in your application system for the integration to properly create these sets.

This arc diagram illustrates the example where and operation can be run on line 1 using operator 1, or on line 2 using operator 2.
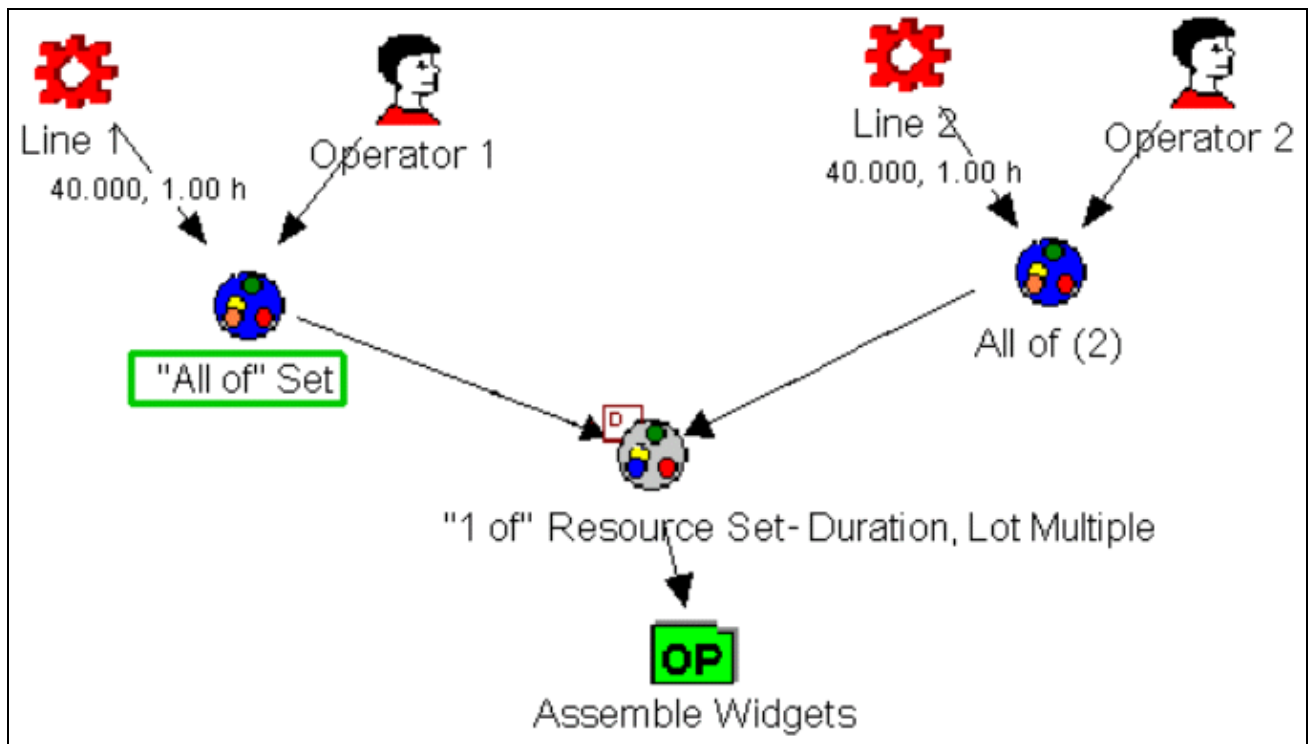
All of set

At the top of the arc diagram, Line 1 and Operator 1 are grouped together as an All of set and Line 2 and Operator 2 are grouped together as an All of set. Line 1 and Operator 1 are the preferred set, which is designated by the green box around the text *All of Set.* Both All of sets point towards the 1 of Resource Set, which means that only one of the All of set can be selected for the operation.

Please note that:

- Any number of resources, such as one or more, can be included in an "All of" Set. Therefore, it is acceptable to have one resource in an "All of" set.

- When an "All of" set is used, the preference is specified on the arc from the "1 of" set to the "All of" set.

- Any crew, machine, or tool can occur more than one time across a structure, provided that it is included in an "All of" set.

- When consistent resource assignment is enabled for the routing, the same resource set is used in subsequent operation steps.

  See "Setting Routing Properties," Creating Production Scheduling Models, Defining Routings, *Production Scheduling*

To create All of sets:

1. Click the Insert "All of" Set button from the toolbar.

2. Move the "All of" Set into the diagram.

3. Click the Insert button on the toolbar.

4. Click the first resource that is to be included in the "All of" set, and then click the "All of" set icon.

   Click the resource and All of set icons in the source-to-destination order. The resource is now linked to the set, as indicated by a connector line with an arrow pointing to the set.

5. Determine the duration and lot multiple for the resource as it relates to the operation.

   Only one resource of each "All of" set can have the duration and lot multiple defined.

6. Repeat steps three and four for the remaining resources for the "All of" set.

   Note. The duration and lot multiple fields are unavailable for additional resources.

   Resources can occur in more than one All of Set.

   ---

7. Repeat steps one through six to create additional "All of" sets for the operation.

   The first set created reads *"All of" Set.* The second set created reads *All of Set (2),* and so one. These sets can be renamed, but must remain unique to each other.

   Note. The All of Set duration and lot multiple can be different across All of Sets.

   ---

8. To link the "All of" sets to the "1 of" Resource Set, click the Insert button on the toolbar.

9. Click the first "All of" set icon, and then click the "1 of" Resource Set icon.

    Click the "All of" set and "1 of" Resource Set icons in the source-to-destination order. The "All of" set is now linked to the resource set, as indicated by a connector line with an arrow pointing to the set.

    When you create the arc between the sets, a Link Properties window appears where you can set a preference for the "All of" set. Enter a 1 in the Preference field to indicate that the set is the preferred set. Enter a 2 in the Preference field to indicate that the set is the second preferred set, and so on for the remaining "All of" sets. The preferred set is indicated with a green box around the "All of" Set text. More than one All of Set can be a preferred set indicated by a 1 in the Preference field.

10. Repeat steps eight and nine to link the remaining "All of" sets to the "1 of" Resource Set.

Changing Multiple Operations

In Production Scheduling, you can change multiple operations at one time to change operation durations and operation lot multiples.

To change the properties of multiple operations:

1. Access the Model Workspace.

2. In the Model Workspace, select the operations that you want to change by holding down the Shift key and highlighting the appropriate operations or folders. If the operations or folders are not listed sequentially, hold down the CTRL key and select the individual operations.

    Note. If you select a folder, all operations in that folder are considered.

3. Right-click and choose Change from the menu.

4. Select the General tab and adjust the necessary duration or lot multiples.

    •   To modify the Discrete Duration, select the appropriate option as needed.

    •   To modify the Discrete Lot Multiple, select the appropriate option as needed.

    Adjusting the Discrete Duration or Discrete Lot Multiples updates the corresponding settings in the Attribute tab of the Link window for the selected operations.

5. Select the Attribute tab and adjust the attribute-base resources assigned the operations. Select the check box to the left of the Name field and, if necessary, assign the appropriate attribute value from the Value list.

6. Click OK to save the changes and close the window.

7. To view the results of the changes, click on one of the modified operations.

Deleting Multiple Operations

To delete multiple operations:

1. Access the Model Workspace.

2. Select the operations that you want to delete by holding down the Shift key and highlighting the appropriate operations or folders. If the operations or folders are not sequential in the list of demands, hold down the CTRL key and select the appropriate operations.

3. Right-click and select Delete from the menu.

Note. You can also multi-select operation shortcuts in the routing diagrams.

Defining Routings

This section provides overviews of routings and discusses how to:

• Work with routing diagrams.

• Create a new routing and routing diagram.

• Set routing properties.

• Work with buffers.

Understanding Routings

A routing is a sequence of operations that are linked together to represent the production of an item. Instead of each product being composed of many different variables that reside in inventory when complete, this approach provides you with a top-down overview of the process. To minimize inventory, when an order for a finished good is received, a routing or series of routings is scheduled to produce all of the items that are required.

Example of a Basic Routing

In this example, a bicycle is manufactured. To complete the bike, several processes must occur, among them: cutting the pieces for the frame, milling the pieces, welding the frame together, painting the bike frame, and final assembly. It is also important that some operations are completed before the next operation can start, that is, Starts After End.

This diagram illustrates a basic process:
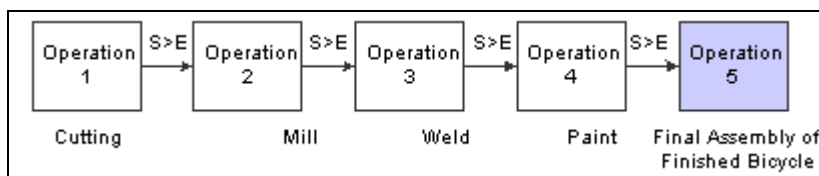
Basic routing of producing a bicycle

In the last operation, you can select to produce a subassembly product, which would then be recorded in inventory and may be consumed by a sales order or by another operation.

Working with Routing Diagrams

This section provides an overview of routing diagrams and precedence relationships, and discusses how to designate primary operations.

Understanding Routing Diagrams

The routing diagram is a graphic representation of the relationships between the operations that make up a routing. Like operations, routings can be organized in folders, if desired. Routings can be oriented in a number of directions for optimal presentation by applying the layout algorithm when you are working in the diagrams.

Within a routing, one operation can be designated as the primary operation. This information is important if you are planning to use the pegging capabilities of Production Scheduling. If you do not select a primary operation in a routing, then you cannot peg "upstream" from the final operation in the routing.

Pegging data is essential for several areas of the application, including various Gantt views, and Key Performance Indicators. To properly generate pegging data, you need to set primary operations in routings and primary output items in operations. The primary operation should be defined as the operation that produces either a subassembly item or the finished product.

The system uses a blue indicator to indicate primary operations in operation diagrams and output items in routing diagrams. You can right-click the operation or item in a diagram and deselect the primary operation option.

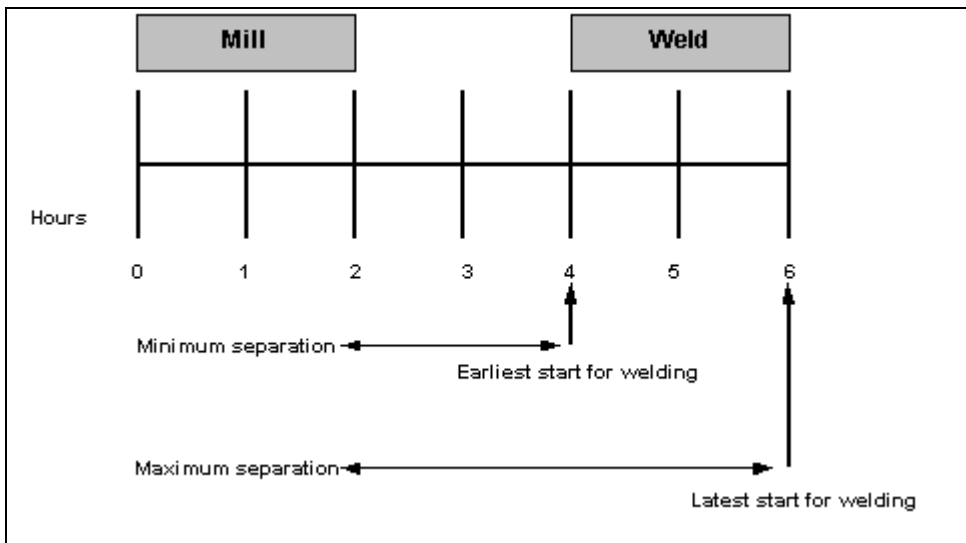### *Understanding Precedence Relationships*

The relationships between the various operations in a routing are specified by the links between operations. These are called precedence relationships, and refer to the relationships and time between operations.

You can specify whether the relationship between two operations is sequential or overlapping. In addition, the amount of time required to move, queue, and set up the next operation is recorded as the minimum separation between the two operations. You can also specify a maximum separation, if required.

In the bike example, if a bike frame must be cut before it is milled, then the relationship of the weld operation to the mill operation is that milling must start after the cutting has finished. The corresponding precedence constraint properties, resulting from a right-click on the box in the connecting line, indicate that the Mill operation is a Starts After End type of constraint. This type of constraint allows for separation times, which, in this example, are a two-hour minimum separation and a four-hour maximum separation.

Using the diagram below as an additional example, assume that the weld operation starts after the two-hour mill operation has finished and the parts have been cooled for another two hours and the welding must start after a maximum of four hours. This dictates that the weld operation must start after two hours of the end of the mill operation.

This illustration represents the precedence relationship:

A typical precedence relationship

Production Scheduling supports the following precedence relationships:

Name	Diagram Label	Description
Starts After End	S>E	One operation starts after the end of the previous operation, according to the time intervals displayed in the Minimum and Maximum separation fields.
Starts At End	S@E	One operation starts at the exact end of the previous operation.
Starts After Start	S>S	One operation starts after the start of the previous operation, according to the time intervals displayed in the Minimum and Maximum separation fields.
Starts At Start	S@S	One operation starts at the same time as the other operation.
Ends At End	E@E	One operation ends at the same time as the end of the other operation.
Minimum or Maximum	+	When S>S or S>E, indicates that the precedence relationship has a minimum or maximum separation

In a routing diagram, links between operations display the precedence relationship as text labels on the links between operations.

The following screen shows the properties for a routing link in which you set precedence relationships:

Precedence Relationship property screen

This diagram illustrates the temporal constraints between operations:
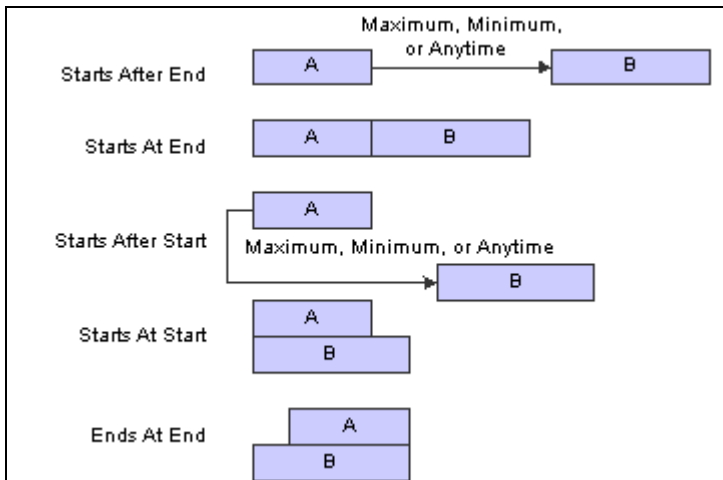
Temporal constraints between operations

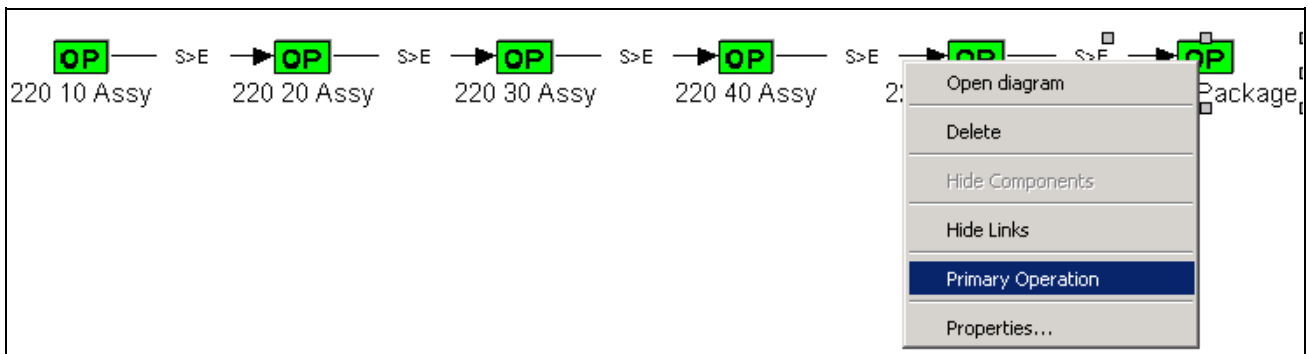### *Designating Primary Operations*

To designate primary operations:

1. Access the Routing Diagram window.

2. Right-click the operation that you want to designate.
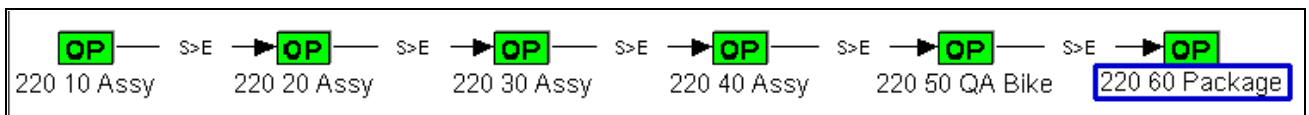
3. Select Primary Operation.

    The operation now appears with a heavy blue border around its name.

    The following diagrams shows the drop down menu, which lets you assign a primary operation.

Designating a primary operation

As a result of selecting Primary Operation, the screen below shows a routing with the last operation designated as the primary operation:

Primary operation

Creating a New Routing and Routing Diagram

To create a new routing and routing diagram:

1. Access the Model workspace:

2. Right-click All Routings or right-click a routing folder and choose New Routing.

3. Type a name for the routing and press Enter.

4. Double-click on the new routing.

    An empty work pane opens.

5. Open the All Operations folder and drag the operations you want to include in the diagram into the empty work pane.

6. On the toolbar, click the Insert Link icon.

7. To link two operations, click the first operation, and then click the second.

Setting Routing Properties

Use this navigation path to access the routing properties window:
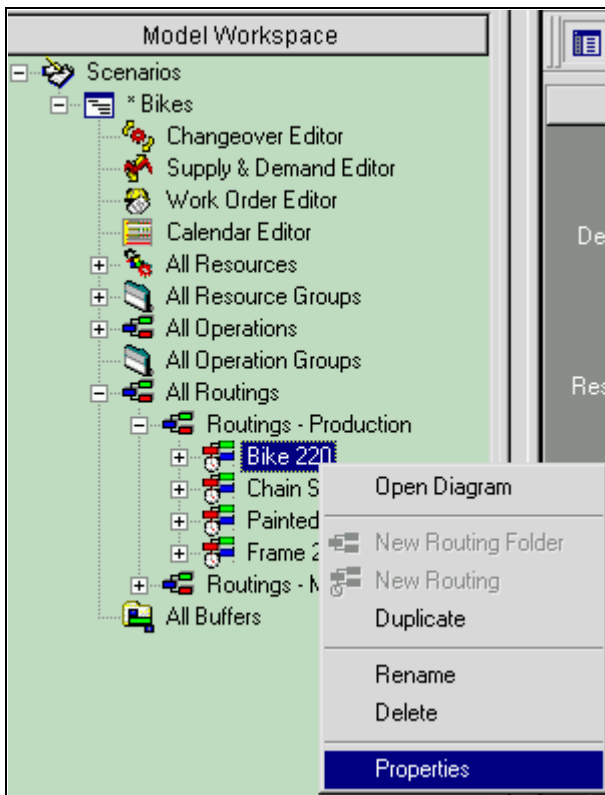
Accessing the routing Properties window

To access the routing properties window:

1. Right-click on a routing.

2. Select the Properties option.

Setting Routing Properties - General

To set general properties complete these fields:

1. From the Model Workspace, right-click on a routing.

2. Select the Properties option.

3. Complete these fields:

Code	Enter a unique code for the routing.
Notes	Enter any routing-specific notes.

Setting Routing Properties - Attributes

To set routing property attributes:

1. From the Model Workspace, right-click on a routing.

2. Select the Properties option.

3. Select the Attributes tab.

4. Complete these fields:

Routing Type	Select a routing type from the list. Options are: • *Production.* Select if your routing produces one or more items. • *Maintenance.* Select if your routing produces no items. This type of routing consumes time and may consume items. For example, a work order to clean a machine would be a maintenance routing. • *Procurement.* Select if your model procures items from external sources. **Note.** If a routing is currently assigned to a work order, you cannot change the routing type or structure without first removing the associated work orders.
Use this Routing for Work Orders only	Select this option if routing should only run when it is associated with a work order. For special configurations or one-off production, you may want to create an alternate routing that produces an item using a different set of operations. The system identifies these routings so that the solver does not try to use them to fulfill demand. Routings identified for use with work orders are only run when they are specified in a work order. If demand exists for an item, but no associated work order exists, the solver does not run a routing that is marked to be used in work orders only.
Schedule Work Order operations according to their unit of effort	This option breaks up associated work order operations into their individual units of effort. By dividing operations into smaller units of effort, scheduling is both more flexible and accurate. This principle is explained in more detail in the examples, which follow this section. See "Understanding Scheduling Work Orders According to Units of Effort," Configuring Solver Options, Setting Up General Solver Options, *Production Scheduling*
Adjacent operations preferred	When enabled, this flag instructs the Production Scheduling solver to schedule all work order operations adjacent to each other when the operations are broken down into their units of effort.
Ship Buffer	Assign a ship buffer if you have enabled Theory of Constraints. See "Working with Buffers," Creating Production Scheduling Models, Defining Routings, *Production Scheduling*

Setting Routing Properties - Resource Assignment

Alternate resources are common in many manufacturing operations. These alternate resources are modeled through the use of resource sets. In many cases, this resource set recurs in an upstream or downstream operation. In many cases, it is important that once a resource is selected in an upstream operation, it must also be selected in a downstream operation.

To perform consistent resource assignment within only a selected group of operations within a routing, setup operation resource assignment groups within a given routing. For example, operation steps 10, 20, and 30 must run on the same resource. However, operation steps 40, 50, and 60, even though they share a common resource with the preceding operations, can run on the alternate.

Use the Routing Properties - Resource Assignment page to set up alternate resources.

To set routing property resource assignments:

1. From the Model Workspace, right-click on a routing.

2. Select the Properties option.

3. Select the Resource Assignment tab.

   All ungrouped operations are listed in the left column and sorted in alphanumeric order.

4. Select the Add New Group button.

   A default group is created.

5. Rename the group by selecting F2 or double-clicking on the group name itself.

6. On the left side of the list, highlight the operation, or operations, that you wish to group together.

7. Click the right arrow button to move the highlighted operations to the right side of the list.

   The operations are assigned to the highlighted group in the right column.

8. Repeat steps two through five to create additional resource assignment groups for any operation within a routing.

Note. If an operation is not assigned to a group it can run on any resource.

If using All of Sets, for example more than one resource such as a machine 1 and tool 1, or machine 2 and tool 2, consistent resource assignment applies to the group. Meaning the group must be the same.

See "Creating All of Sets," Creating Production Scheduling Models, Defining Operations, *Production Scheduling*

Setting Routing Properties - Where Used

The Routing Properties - Where Used page displays:

• Routing in the event that a routing is embedded in another routing.

• Work orders that contain this routing.

You can double-click a routing or a work order in the list to open it.

Working with Buffers

Production Scheduling provides you with the option to apply Theory of Constraints principles and applicable business processes to your scheduling environment. Theory of Constraints principles can be enabled in a particular schedule in order to take advantage of Buffer Planning functionality.

This section discusses:

*   Drum buffer rope

*   Shipping buffers

*   Capacity constrained resource buffers

*   Assembly buffers

*   Supplier buffers

*   Buffer validation rules.

*   Define buffer attributes.

*   Assign buffers to routings.

*   Assign buffers in routing diagrams.

*   Assign a ship buffer to a routing.

Drum Buffer Rope

Production Scheduling incorporates Drum Buffer Rope (DBR) methodology through Buffer Planning functionality. The DBR methodology is designed to balance the flow of materials or work in progress items across the production cycle. In order to achieve this balance, production is synchronized with the capacity constrained resource (CCR) that is considered to be the "weakest link" in the production process. The DBR methodology is based on the principle that production in a supply chain is limited to the throughput on the resource or resources with capacity constraints.

The drum represents the production rate of the CCR. All production in the supply chain is determined by the rhythm of the drum. The rope ensures that materials are not inserted into the operation with the CCR at a faster rate than the CCR can process. Shipping, CCR, assembly, and supplier buffers allow you to build protection mechanisms within your production schedule to combat production variability.

You can use buffers in your model to help you manage the schedule variability that can occur in a manufacturing environment that depends on one or more CCRs. When you define buffers in your model, you specify a target buffer value and a minimum buffer value. The solver target buffer value is a soft constraint. The minimum buffer value is a hard resource constraint. In situations where there is insufficient resource capacity to produce a demand item on time, the minimum buffer value can override the target buffer value and cause a delay in the schedule for a demand item. For all buffer types, the buffer value must be greater than or equal to the sum of the durations in the routing and minimum separation times between operations. Although you can create multiple buffers in a linear routing, only one buffer can exist between a pair of operations within the routing.

This diagram illustrates the use of buffers in a routing:
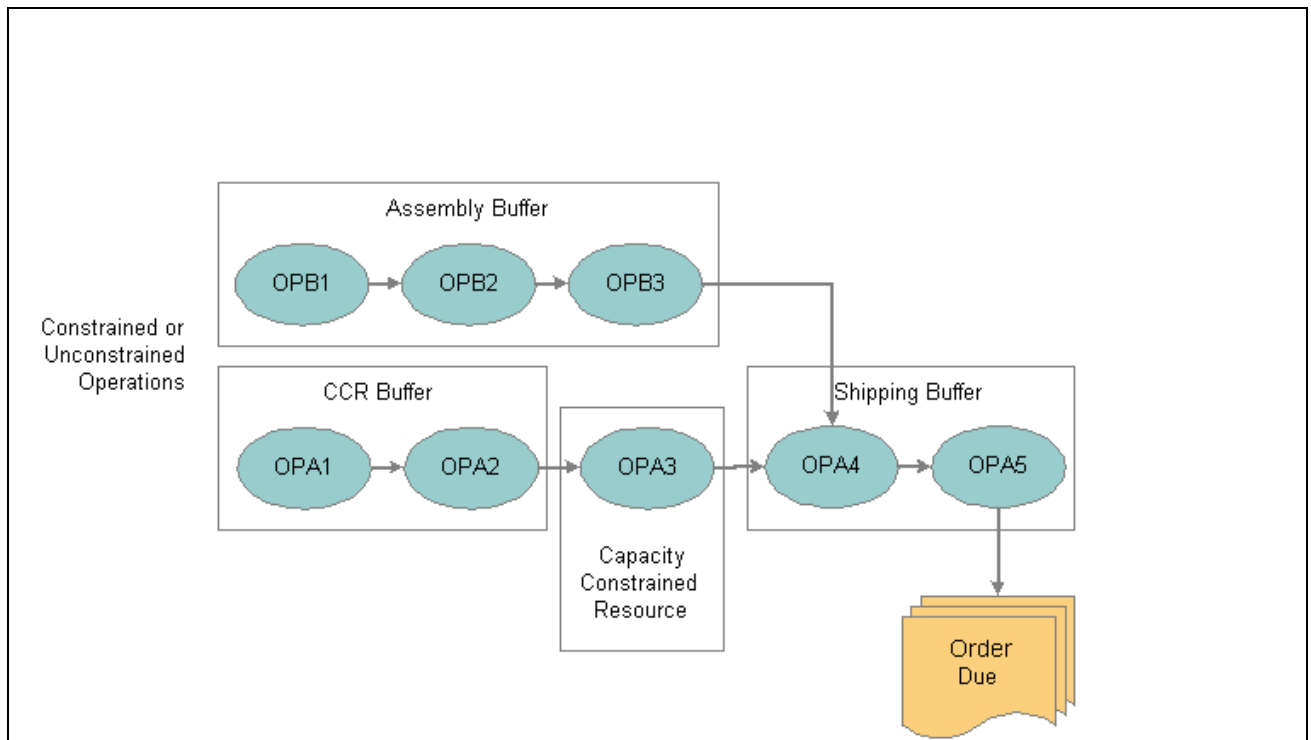
Buffers within a routing

Shipping Buffers

In a lean manufacturing environment that recognizes Theory of Constraints principles, a shipping buffer can be used to estimate the time between the point where an item is processed by a CCR and the point where the item is shipped. The purpose of the shipping buffer is to protect the shipping due date from downstream variability in the supply chain.

It is necessary to define a CCR in a routing for the solver to respect a shipping buffer. In certain scenarios where two operation paths combine to make a routing, the solver calculates the shipping buffer by identifying the path with the shortest duration between the primary operation and the CCR operation. The solver subtracts the shortest duration from the target and minimum buffer sizes and uses this value as the true buffer value.

Capacity Constrained Resource Buffers

Capacity constrained resource buffers (CCR buffers) are intervals of time that precede a CCR in a routing. CCR buffers are used to protect the schedule from variability caused by operations that are downstream from the CCR.

You can incorporate CCR buffers into your model in order to avoid CCR overload and CCR starvation. CCR overload occurs when work in progress items accumulate at the CCR due to a rate of material insertion that exceeds the CCR capacity. CCR starvation occurs when a CCR becomes idle due to lack of work in progress items caused by an upstream disruption in the manufacturing process. CCR starvation is often caused by a rate of material insertion that is sufficiently lower than the CCR capacity, causing the CCR to be "starved" for work.

Within a routing, operations that use a CCR are known as CCR operations. CCR buffers must be defined for CCR operations. In a linear routing, you can define as many CCR buffers as you require to effectively model bottlenecks in your manufacturing process. You must have one or multiple CCRs defined in your model before you begin assigning CCR buffers. Each resource that is designated as a CCR appears highlighted with a red border in the routing diagram.

Assembly Buffers

Assembly buffers are an extension of shipping buffers that you can use in a model to protect a schedule from variability that may be incurred when a non-constrained item is combined with an item that is processed by a CCR. Assembly buffers specify the amount of lead time from the release of materials to an operation that involves assembling a non-CCR processed item with an item processed by a CCR.

The solver respects assembly buffers that are appropriately defined in the model. The solver does not respect assembly buffer values in the following cases:

- The beginning or terminating operation used to define the buffer value is a CCR operation.

- A CCR operation exists between the beginning or terminating operation used to define the buffer value.

Supplier Buffers

You can define a supplier buffer for each item that is procured from a supplier. Supplier buffers are used by Production Scheduling to calculate arrival dates for purchase orders. Supplier buffers represent a period of time between the supply event for an item and the beginning of the operation that uses the item. The supplier buffer enables you to build in a buffer to accommodate time between the arrival of the item from the supplier and the use of the item in an operation. Operations involving items with a supplier buffer begin at the end of the supplier buffer. To use the supplier buffer functionality, the Theory of Constraints must be enabled for the model or for a specific schedule.

A supplier buffer is not an extension of the lead time required to supply the item. You can specify a target buffer value and a minimum buffer value when you define supplier buffers in your model. When you define supplier buffers, the following guidelines must be observed:

- The minimum buffer value must be greater than or equal to the lead time of the purchased item.

- The target buffer value must be greater than or equal to the combined values of the minimum buffer and the lead time between the estimated ship date and the arrival date of the purchased item.

Buffer Validation Rules

Production Scheduling has built in buffer validation rules that govern buffer behavior during a solve. You should observe the following buffer validation rules while you create buffers in a routing:

- Within a resource set, if a CCR is defined, every resource in the set must be defined as a CCR.

- The minimum buffer value must be less than the target buffer value.

- A negative buffer value is not allowed.

- The buffer value in a routing must be greater than the sum of the durations and separation times between connected operations.

- Buffers can not be defined in reverse in a routing.

- A maximum of one buffer is allowed per relationship.

- The solver ignores buffers in any CCR constrained routing that does not have buffers defined in the routing or associated to the routing.

- All operations using a specific CCR should have a CCR buffer defined.

Defining Buffer Attributes

To define buffer attributes:

1. Access the Buffer Properties window.

2. Select the Attributes tab.

3. Complete these fields:

Target Buffer Size	Specify a target buffer value. You can specify a period of time in seconds, minutes, hours, and days. The value can include one decimal point. This value is a soft constraint.
Minimum Buffer Size	Specify a minimum buffer value. You can specify a period of time in seconds, minutes, hours, and days. The value can include one decimal point. This value is a hard constraint.

Assigning Buffers to Routings

To assign buffers to routings:

1. Access the Routing Properties window.

2. Select the Attributes tab.

3. Select a predefined buffer from the available options in the list box.

4. Click Apply.

5. Click OK.

Assigning Buffers in Routing Diagrams

To assign buffers in routing diagrams:

1. Access the Routing Diagram window.

2. Click the Buffer Link button on the toolbar.

3. In the operation diagram, click an operation to select the operation.

4. To create a buffer between two operations, click the first operation, and then click the second operation in the operation diagram.

5. In the Buffer Properties dialog box, select a predefined buffer from the available options in the Buffer list box.

6. Click OK.

Assigning a Ship Buffer to a Routing

To assign a ship buffer to a routing:

1. From the Model Workspace, right-click the routing.

2. Select Properties.

3. Select the General tab.

   Note. In order to assign buffers, the Theory of Constraints option must be enabled on the Scenario Properties - General page.

   See "Setting Up the Scenario Properties - General page," Configuring Scenario Properties, *Production Scheduling*

4. In the Code field, select a pre-defined buffer.

Defining Work Patterns and Availability (Resource Calendars)

This section provides an overview of work patterns and availability and discusses how to:

• Create a global calendar

• Set calendar horizons

• Add calendar events

• Modify calendar events

• Delete calendar events

Understanding Work Patterns and Availability

The Calendar editor enables you to set up work time, such as work patterns and planned down time. These calendars are assigned to resources on the Resource Properties page. You can view calendar event properties to review the start date, end date, and duration. In addition, you can revise the information, but revisions appear only after the schedule has been repaired.

Work patterns specify the type of work operation, such as two shifts, three shifts, or five days a week. Down time includes interruptible (delay) and non-interruptible (down) times. Operations can start before and continue after a delay. Operations cannot start before and continue after down time.

Planned down time periods include vacation shutdowns, planned maintenance, delays, and scheduled shutdown days. For example, if a plant works only during weekdays, then the calendar displays weekends as down time periods. By creating a calendar for the down times of each resource, you tell the system how to plan production more precisely and allocate resources more effectively.

To associate a calendar to a crew, machine, or tool, use the Availability tab on the Resource Properties page.

See "Defining Resource Availability," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

Creating a Global Calendar

To create a global calendar:

1. Access the Calendar editor.

2. Click Add Calendar.

3. Specify a name for the calendar in the Calendar Name field.

4. Click OK.

Setting Calendar Horizons

To set calendar horizons:

1. Access the Global Calendar Horizon window.

2. Complete these fields:

Start Time	Specify the beginning of the calendar horizon. You can enter the date manually in the field or use the calendar tool.
End Time	Specify the end of the calendar horizon. You can enter the date manually in the field or use the calendar tool.
Duration	Specify the duration as a value in seconds, minutes, hours, or days.

Adding Calendar Events

To add calendar events:

1. Access the Add Calendar Events window.

2. Complete these fields:

Uptime	Select this option to indicate that all machines and activities are available.
Downtime	Select this option to indicate down time when all machines and activities are stopped.
Delay Time	Select this option to indicate a delay when machines and activities are temporarily suspended.
Start Time	Specify the date and time when the event begins. Alternately, you can leave the original date. When the end time and the duration are entered, the start time is automatically calculated.
End Time	Specify the date and time when the event ends. Alternately, you can leave the original date. When the start time and the duration are entered, the end time is automatically calculated.
Duration	Specify the event length in seconds, minutes, hours or days. If the start time is specified, the end time is automatically filled in, based on the duration. Similarly, if the end time is specified, the start time is filled in automatically. If both start and end time are specified, the duration is automatically calculated.
Recurring event	Select this option if you want the calendar event to repeat.
Daily	Select this option if you want the calendar event to repeat daily.
Weekly	Select this option if you want the calendar event to repeat weekly.
Monthly	Select this option if you want the calendar event to repeat monthly.
Yearly	Select this option if you want the calendar event to repeat yearly.
No limit	Select this option if the event occurs throughout the calendar horizon.
End after	Select this option, and select the number of occurrences if the event occurs for a specified number of times and then stops.
End by	Select this option, and select an end date if the event ends by a specified date.

Modifying Calendar Events

To modify calendar events:

1. Access the Edit Calendar Event window.

2. Make any modifications to the event type and the event time and duration data.

3. Select one of the following options:

   • Edit Current Selection

      Modify the current calendar event.

   • Edit Matching Events In Range

      Modify all calendar events in the date range specified. Complete the Start Time and End Time fields. The default value of the End Time field is the end of the scheduling horizon.

4. Click OK.

Deleting Calendar Events

To delete calendar events:

1. Access the Calendar editor.

2. Click a calendar event in the list. To select multiple events, press Shift, for contiguous selections, or Ctrl, for non-contiguous selections.

3. Select one of the following options:

   • Remove Event

      The event is removed from the schedule.

   • Remove All Events

      All calendar events in the date range are removed from the schedule.

4. Click Yes to confirm the deletion of the calendar event.

Specifying Supply and Demand Events

This section provides an overview of the supply and demand editor and demand priorities, and discusses how to:

• Create supply and demand folders.

• Create supply orders.

• Create demands.

• Define demand policies.

• Change multiple demands.

• Delete multiple demands.

• Create work orders from demands automatically.

Understanding the Supply and Demand Editor

The Supply & Demand editor enables you to set up supply and demand events for your schedule. Although this data is usually imported into Production Scheduling, you can input supply and demand events directly within the application to facilitate simulations.

You can create demand orders and demand line items by creating a demand folder, adding demand orders, adding line items, specifying demand policies and dates, and completing the line item fields. You can edit demand information in either order view or line item view.

Supply and demand events can be modified globally or for a particular schedule. The global Supply & Demand editor is launched from the Model Workspace, or an icon on the toolbar. The Supply & Demand editor can be modified locally for a specific schedule. This ability enables you to revise requirements or purchased parts that are required for a schedule, and analyze what-if scenarios.

In addition to the maintenance of supply and demand orders and their status, the Supply & Demand editor also contains an analytical tool that provides pegging information for all demand orders. This pegging grid displays all of the items and processes that are involved in the manufacturing of a product; it also displays when each step is scheduled for completion. This grid provides you with specific production information for a particular line item when such information is required.

Understanding Demand Priorities

In most manufacturing environments, demand priority setting is a common practice. These priorities are generally dictated by various factors, such as customer size, revenue contribution, length of time a sales order has been outstanding, predetermined agreements, as well as other factors. In the event of resource conflicts, higher priority orders are expected to receive preferential treatment over other orders in the system.

Demand priorities often change during planning simulations, such as when a scheduler is doing ad-hoc order promising to facilitate the effect of adding a hot order to the existing production schedule, that is, in which other demands are late or otherwise affected. This can be an iterative process to achieve the schedule and service levels that satisfy the user.

Production Scheduling provides very differentiating scenario comparison capabilities to support these types of scenarios.

In the Supply & Demand Editor, there are three order classifications:

• Hot Order

• Committed

• Uncommitted

Within each classification, you can set a demand priority field from one to five. Therefore, in effect you have fifteen different priorities with which to set your demand. The lower the number in the Demand Priority field, the higher priority that a demand item is given. For example, the solver fills demand orders with a priority of 1 before filling orders with a priority of 2.

If two or more demand orders have the same requested date, then the system considers demand priority for each order. The solver schedules the demand that has the lower priority.

The highest order class is Hot Order, followed by Committed, and then Uncommitted. The highest Demand Priority is 1, and the lowest Demand Priority is 5. Uncommitted orders are essentially opportunistic orders which look for holes in the schedule and inventory to fill a given demand.

You set the class and priority in either the Rules tab of the order, or in the line item itself.

The Production Scheduling solver schedules higher priority demands first in the event of a resource conflict. If no resource conflict exists, the schedule is driven by due date, which enables the lower priority orders to be scheduled first.

Example of Demand Priority

The importance of supporting these scenarios, and the shortcomings that result when these types of scenarios are not supported is best demonstrated in the following example.

Two demands, Demand 1 and Demand 2, have stated priorities of 2 and 1 respectively. Demand 1 is due on Day 1 and Demand 2 is due on Day 2. Each demand shares a common machine, which is used to manufacture the respective items on the demand.

If there are no resource conflicts, that is, there is sufficient capacity, demand 1 is scheduled first because the lower priority demand is due earlier. Demand 2 is scheduled next because the higher priority demand is due after Demand 1.

If these demands compete for capacity, users expect that the higher priority demand would be fulfilled first.

Creating Supply and Demand Folders

To create supply and demand folders:

1. Access the Supply & Demand editor.

2. Right-click All Demands or All Supplies, and select Add Folder.

3. Right-click the folder, select Rename, and enter a new name for the folder.

4. Press Enter.

Creating Supply Orders

To create supply orders:

1. Access the Supply & Demand editor.

2. Right-click the Supply folder to which the new supply order is to be added.

3. Select Add Supply to add an order.

   An order appears under the folder with a system-generated number.

4. On the General tab, complete the following fields:

Supply Number	The supply order is automatically assigned a supply number. You can edit this value if necessary.
Origin Address	Specify the address from which the vendor is shipping the supplies. This field is optional.
Type	Select *Allocated Supply,Interplant Shipment,* or *Purchase Order.*
Status	Select *Approved* for the solver to consider the supply order, or *Hold* for the solver to ignore the supply order.

5. Right-click the order in the tree and select Add Line Item.

6. Complete the following line item fields:

Item	Select the item from the drop-down list. Manufactured and purchased items are displayed.
Quantity	Specify the quantity being supplied.
Status	Select the approval level of items supplied: *Approved* or *Hold.*
Earliest Date	Optionally, specify the earliest date when the order can be delivered. The default is the start of the horizon. If the date that is entered precedes the schedule horizon start, then this date is changed to the horizon start when solving. If the date entered follows the schedule horizon end, then the solver ignores this order.
Requested Date	Specify a date when the order is expected to arrive. The solver attempts to keep the supply event as close to the due date as possible. The solver is free to select a date within the earliest and latest date range if required.
Latest Date	Optionally, specify the latest date when the order must be delivered. The default is the end of the horizon. If the date that is entered falls after the schedule horizon end, then this date is changed to the horizon end when solving. If the date entered precedes the schedule horizon start, then the solver ignores this order.

7. Click Save.

Note. Available Date is a system-generated, display-only field.

You must have created or imported item data before a line item can be added.

Creating Demands

To create demand orders:

1. Access the Supply & Demand editor.

2. Right-click the Demand folder where you want to add the order and select Add Demand.

    An order appears under the folder with a system-generated number.

3. On the General tab, complete the following fields:

Demand Number	The demand order is automatically assigned a demand number. This value can be edited, if necessary.
Customer Name	This field is optional.
Customer Number	This field is optional.
Ship To Address	This field is optional.
Type	Select *Order, Safety,* or *Forecast.*
Status	Select *Approved* for the solver to consider the demand order or *Hold* for the solver to ignore the demand order.

4. Click Save.

Defining Demand Policies

Use this page to define demand policies:

Defining demand policies

To define demand policies:

1. Access the Supply & Demand editor.

2. Select a demand order from the tree.

3. On the Rules tab, complete the following fields in the Demand Policies section:

Order class	Select *Hot Order,Committed,* or *Uncommitted.*
Backorders Allowed	Select *Yes* to enable the requested date to be violated up to the latest order. Select *No* if all line items must be available at the requested date.
Demand Priority	Choose a value from 1 to 5.
Ship Complete	Select *Yes* to ship the order only when all of the line items have been produced. Select *No* to enable partial shipments of the order.

Note. Available Date is a system-generated, display-only field.

4. To add a new line item to the order, right-click the order in the tree and select Add Line Item.

5.   For the line item, complete the following fields:

Item	Click the arrow by the code for the item and choose from the available options. Only saleable items are displayed.
Quantity	Specify the quantity of the item that is requested.
UOM (unit of measure)	Displays the unit of measure for the tiem.
Status	Select the approval level of the item that is requested. The status can be *Approved* or *Hold.*
Type	Select one of the following options:

- *Forecast,* which indicates that the aggregated demand estimates for a specific time period is based on historical sales information.

- *Orders,* which indicates that the actual orders placed by customers. Safety Stock. Items that are needed to maintain safety stock levels.

Demand Priority	Specify a relative weighting (degree of build) with 1 being the highest priority. For example, with two JIT orders, the order with demand priority 1 is produced as late as possible, relative to the order with demand priority 2.
Requested Date	Specify the requested date for the item that is requested.
Earliest Date	Optionally, specify the earliest date when the order can be delivered. The default is the start of the horizon. If the date entered precedes the schedule horizon start, then this date is changed to the horizon start when solving. If the date entered follows the schedule horizon end, then the solver ignores this order.
Latest Date	Optionally, specify the latest date when the order can be delivered if backordered. The default is the end of the horizon. If the date entered falls after the schedule horizon end, then this date is changed to the horizon end when solving. If the date entered precedes the schedule horizon start, then the solver ignores this order.

6.   Click Save.

Changing Multiple Demands

In the Supply and Demand Editor, you can change the properties of multiple demands at one time.

To change the properties of multiple demands:

1.   Access the Supply & Demand editor.

2. Select the demands that you want to change by holding down the Shift key and highlighting the appropriate demands, folders, or line items. If the demands, folders, or line items are not listed sequentially, hold down the CTRL key and select the individual demands.

    Note. If you select a folder, all demands that are in that folder are considered.

3. Right-click and choose Change from the menu.

The menu is shown in the following screen:
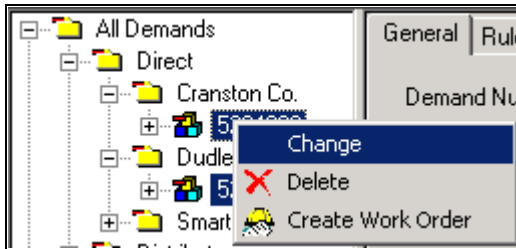
Change demand option

Changing General Properties of Demands

Use General page of the Change — Demand window to change demand policies and demand priorities.

1. In the General window, select the policies and priorities you want to change by clicking the box next to the policy or priority type.

2. From the drop-down menus, beside the policies and priorities, select an option.

3. Click the Details tab and click the box beside the property you want to change.

4. From the drop-down menus, beside each property detail, select an option.

    Note. An explanation of each field and its possible values is provided immediately following this procedure.

5. Click OK to save the changes and close the window.

6. To view the results of the changes, click on one of the modified demands.

Field definitions are:

Demand Type	Choose Order, Forecast, or Safety Stock.
Status	Choose *Approve* or *Hold.*
Backorders Allowed	Choose No or Yes.
Ship Complete	Choose No or Yes.
Order Class	Choose Hot Order, Committed, or Uncommitted.

Priority	Choose 1, 2, 3, 4, or 5. One is the highest priority and five is the lowest priority.

Changing Detailed Properties of Demands

Select the Details tab in the Change ― Demand window. Use this page to change demand quantities, request dates, earliest dates, and latest dates.

The properties and the options that are available from the drop-down menu in the Details page are listed below.

In the Demand Quantity region:

Do Not Change	Leave demand unchanged.
Set to	If you want all your demands to be the same quantity, enter that quantity in this field.
Multiply by	Enter the factor by which you want to multiply the demands. For example, to double demands, enter 2.00. To half the demands, multiply by 0.5.

In the Request Date region:

Do Not Change	Leave request date unchanged.
Set to	To set all the request dates of the selected demands to the same date, click the up or down arrow in this field to move the date forward or backward, respectively.
Advance by	You can bring the request date forward (closer to the current date), or move it back (farther from the current date) by using this field. To move the request dates forward, enter the number by which you want to advance it in the Advance by field. Select the increment of time from the drop-down menu. Valid values are: *minute,hour,day,* or *week.* For example, to advance the request date by 10 hours, enter 10 in the Advance by field and choose hour from the drop-down menu.
Reduce by	You can move the request date back (farther from the current date) by using this field. To move the request dates back, enter the number by which you want to reduce it in this field. For example, to move the request date back by 10 hours, enter 10 in the Reduce by field and select hour from the drop-down menu. Valid values are: *minute,hour,day,* or *week.*

In the Earliest Date region:

Do Not Change	Leave earliest date unchanged.
Set to blank	Select this option to change all earliest dates to empty fields, which are optional.

Set to	To set all the earliest dates of the selected demands to the same date, click the up or down arrow in this field to move the date forward or backward, respectively.
Advance by	You can bring the earliest date forward (closer to the current date) by using this field. To move the earliest dates forward, enter the number by which you want to advance it in the Advance by field. Select the increment of time from the drop-down menu. Valid values are: *minute,hour,day,* or *week.* For example, to advance the earliest date by 10 hours, enter 10 in the Advance by field and choose hour from the drop-down menu.
Reduce by	You can move the earliest date back (farther from the current date) by using this field. To move the earliest dates back, enter the number by which you want to reduce it in this field. For example, to move the earliest date back by 10 hours, enter 10 in the Reduce by field and select hour from the drop-down menu. Valid values are: *minute,hour,day,* or *week.*

In the Latest Date region:

Do Not Change	Leave latest date unchanged.
Set to blank	Select this option to change all latest dates to empty fields, which are optional.
Set to	To set all the latest dates of the selected demands to the same date, click the up or down arrow in this field to move the date forward or backward, respectively.
Advance by	You can bring the latest date forward (closer to the current date) by using this field. To move the latest dates forward, enter the number by which you want to advance it in the Advance by field. Select the increment of time from the drop-down menu. Valid values are: *minute,hour,day,* or *week.* For example, to advance the latest date by 10 hours, enter 10 in the Advance by field and choose hour from the drop-down menu.
Reduce by	You can move the latest date back (farther from the current date) by using this field. To move the latest dates back, enter the number by which you want to reduce it in this field. For example, to move the request date back by 10 hours, enter 10 in the Reduce by field and select hour from the drop-down menu. Valid values are: *minute,hour,day,* or *week.*

Deleting Multiple Demands

In the Supply and Demand Editor, you can delete multiple demands at one time.

To delete multiple demands:

1. Access the Supply and Demand editor.

2. Select the demands that you want to delete by holding down the Shift key and highlight the appropriate demands, folders, or line items. If the demands are not sequential in the list of demands, hold down the CTRL key and select the appropriate demands.

3. Right-click and select Delete from the drop-down menu.

Creating Work Orders from Demands Automatically

To automatically create multiple work orders that correspond to one demand line:

1. From the Model Workspace, open the Supply and Demand Editor.

2. Right-click on one demand line item, one folder, or use the multi-select option, and select Create Work Order.

    The application automatically creates work orders that map one-to-one for each demand that you selected.

 The folder structure from the Supply and Demand Editor is duplicated for each work order. If a demand can be fulfilled from multiple routings, the application prompts you as to which routing the work order should use. All quantities and dates on the work order are inherited from the demand.

If a work order is already linked to a demand, a warning dialog appears and prompts the user if they wishes to create another work order for that demand. If alternate routings exist, a dialog box appears and prompts the user to select a routing.

Working with Work Orders

This section provides overviews about work orders, flexible scheduling of work order, automatically generated routings base on work orders, and work order yield and scrap, and discusses how to:

• Create work order.

• Delete a work order.

• Create work order folders.

• Organize work orders into folders.

• Create hard links between work orders.

• Release work orders to production.

Understanding Work Orders

Work orders are statements of current or future production for an item. In Production Scheduling, work orders can be linked to demand orders. Work orders are usually imported from a Materials-Requirement Planning (MRP) system, although they can also be entered directly into the application. Production Scheduling supports three types of work orders:

• Production orders

- Maintenance orders

- Engineering changes

A production work order must specify an item that is produced on completion of all of the operations within the work order. Only manufactured items (as defined in the item properties) are available for selection. You then select a routing that produces the item.

Maintenance orders and engineering changes are a secondary type of work order. Maintenance orders and engineering changes do not produce items, but they still consume production capacity. These orders play a crucial role in the manufacturing process. For example, machines break down if they are not properly maintained, and machines and equipment are subject to enhancement due to technological and engineering advances. Maintenance orders support and repair anything involved in the production process. For example, maintenance orders can be issued to repair a drill press, inspect welding equipment, or lubricate a milling machine. xx

Engineering changes are used when machines and tools require capital improvements to perform new or improved operations.

Although work orders that are listed in the Work Order Editor and demand that is displayed in the Supply & Demand editor are different sources of production requests, they can be linked. The system considers both requirements when the model is scheduled.

Work orders can also be configured individually when the work order is not in the standard template when it is imported from an ERP system.

Working with Maintenance Work Orders

Maintenance work orders in the context of Maintenance activities represent specific tasks to be completed during a maintenance visit. These tasks typically define material and resource requirements and are grouped into logical Stages which are sequential blocks of tasks within a Visit. The completion of a stage represents a Milestone and must be completed before the next Stage can begin. In many cases, there is a sequence of tasks that must be followed and can be serial or parallel in nature. This complex relationship is referred to as the task hierarchy which needs to be respected when defining a maintenance schedule. Maintenance work orders contain this necessary information so that they can be scheduled appropriately.

One of the key differences of a maintenance work order compared to a production work order is that it may not necessarily produce an item to stock at the completion of the single work order. Recall that a series of maintenance tasks are often required and it may not be until the last task, or work order, which produces a unit to stock. For example, an engine may require a series of maintenance tasks and it is not until the last work order that the engine can be deemed usable, or produced to stock. Another fundamental difference of a Maintenance work order in the context of a visit is the Work Order Request date does not guide scheduling of the work order. Instead, the Visit request date is used to drive the start and completion date of the work order considering the various resource and material constraints in addition to the complex task hierarchy and Milestones constraints that are relevant to the Visit. The Visit Request date governs the work order scheduling within the Production Scheduling Solver.

Flexible Scheduling of Work Orders

This section discusses:

- Scheduling work orders inside the fixed time fence.

• Scheduling work order outside the fixed time fence.

Scheduling Work Orders Inside the Fixed Time Fence

In Production Scheduling, you can set a fixed production timefence on a given resource. This fixed production timefence dictates that any new production that is required to fulfill demand must occur beyond the end of the timefence duration.

Within a fixed production timefence, you can have work orders that are specified to run, which are created either from a previous planning cycle or by another process. These work orders may have operations with a fixed earliest start date, a fixed latest end date, or both. Because work order operations within a fixed timefence are considered fixed, a situation may arise where an infeasibility occurs within the solve process. For example, if there are two operations on different work orders that run on the same machine that have overlapping fixed dates, an infeasible solution may occur within the solve process. This diagram illustrates this concept:
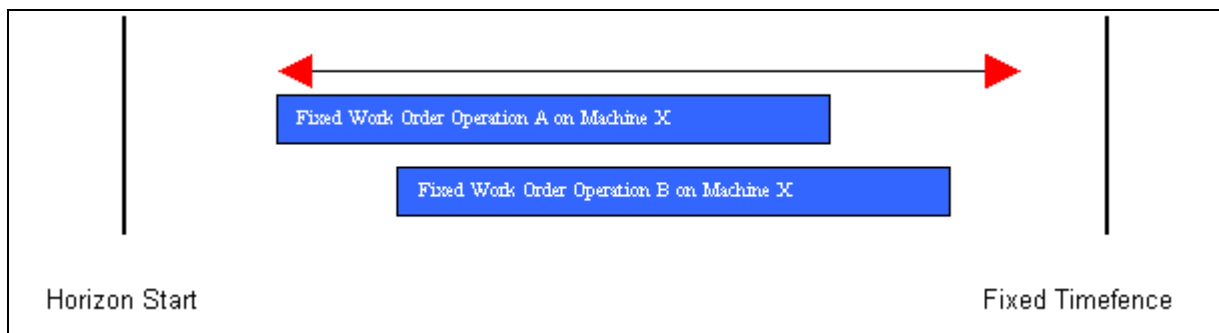
Two different work orders on the same machine

Note. To fix operations on a work order within the fixed timefence, you must populate the Earliest Start Date field, or the Latest End Date field, or both. The fix dates must dictate that all or some of the operation duration falls within the resources fixed timefence.

To address the possibility of an infeasibility within the fixed timefence, the Production Scheduling solver may move work order operations around and possibly schedule work order operations to occur prior to the start of the planning horizon. This indicates that your shop floor is running behind.

If any open work order operations, within the fixed timefence, are moved to start within the timefence or before the horizon start to avoid an infeasibility, you receive either a Work Order Past Due or Work Order Fixed Date Violation alert. The alert tells you the work order this situation has occurred, the operation, and the relevant start and end times.

If you receive a Work Order Past Due alert in order to avoid an infeasibility, this indicates that a work order operation has been moved within the timefence and is scheduled to have a start time that occurs prior to the horizon start. The violation start-time is the time the operation starts within the past. The violation end time is the end of the operation.

If you receive a Work Order Fixed Date Violation, this indicates that a work order operation has been moved within the fixed timefence and is scheduled to have a start time different from the start date that is specified on the work order routing. The violation start time is the time the operation starts. The violation end time is the end of the operation.

You can filter these alerts using the filtering functionality in the Production Scheduling alert window to be one of two types: Work Order Operation Past Due or Work Order Fixed Date Violation. By performing a right-click on the alert in question, you can navigate to the Work Order Editor routing tab to further investigate.

When scheduling work orders in conjunction with a fixed timefence on a resource, the solver behaves in one of the following ways, depending on the circumstances:

1. When work order operations do not have a start date or a latest end date specified, that is, the Earliest Start Date or Latest End Date fields are blank:

    •   If a work order request date would suggest a work order be scheduled in the timefence, it is scheduled beyond the timefence.

    •   If a work order is attached to a demand and the demand has a latest date that would suggest the work order be scheduled in the timefence, the work order can be scheduled in the timefence.

2. When work order operations have a start date, that is, the Earliest Start Date field is completed.

    •   All work order operations ending before the time fence are considered "unfixed" in the solver, so that they can move within the timefence or before the timefence.

    •   All work order operations that cross the time fence (start before and finish after the time fence) are fixed to their start times, unless an infeasibility occurs.

    •   If the total capacity of all the fixed operations that cross the time fence boundary exceed the machine's capacity, the latest operation is fixed and the others are moved earlier.

    •   If there is not enough on-hand inventory, that is, starting inventory, work orders are not moved before the horizon start and are pushed to start after the fixed timefence or the first inventory injection based on supplier lead time or a supply event (which can be another work order producing the material or supply).

    •   If a work order operation is fixed to start before the horizon start, it is scheduled at that time.

3. When work order operations have a latest end date field, that is, the Latest End Date on the work order operation is completed.

4. All work order operations ending before the time fence are considered unfixed in the solver, so that they can move, if necessary, within the timefence or before the timefence.

    •   All work order operations that cross the time fence (start before and finish after the time fence) are fixed to their end times.

    •   If the total capacity of all these operations that cross the time fence boundary exceed the machine's capacity, based on priority and due date and the earliest start date, the conflicting operations are moved to start before the timefence, when feasible, and not pushed later.

    •   If there is not enough on-hand inventory, that is, starting inventory, work orders are not moved before horizon start and if they can not be moved past a given end date, then the application informs the user that there is not enough inventory, that is, no producing operation error message.

    •   If a work order operation has a latest end before the horizon start, it should be scheduled at that time if possible.

5. Active Work Order Operations

    •   If work order operations are *Active,* they are generally scheduled at horizon start even with a fixed timefence. However, if there is no time fence, and two work order operations are marked as active on a single capacity resource, then one of them is moved to avoid an infeasibility.

    •   If there is a time fence, then all active work order operations are unfixed and moved around according to due date. If this action pushes other operations out beyond the fixed timefence, then the solver moves it earlier. In this case, the work order attached to the lower priority demand is moved. If there are no demands or the priorities are the same, one of them is moved based on request date, then on alphanumeric key of the order itself.

    •   If a work order operation is marked as *Active* and is also considered to be fixed, the fix date is ignored.

Scheduling Work Orders Outside the Fixed Time Fence

The solver resolves conflicts where work orders are outside of the fixed timefence to avoid situations such as:

•   More than one work order operation is firm on the same single capacity resource and they have overlapping date and times.

    If this occurs, the work order with the highest priority, followed by earliest request date, is scheduled first. A supporting alert is provided to indicate a firm date violation.

•   A firm date range on a work order operation is shorter than the total operation run and elapsed time.

    For example, if the firm range is 1.5 hours and the operation run time is 9 hours. If this occurs, the work order operation is still scheduled. A supporting alert is provided to indicate an operation firm date violation.

Automatically Generated Routings Based on Work Orders

Manufacturers generally fall into one of two categories, make-to-order (MTO) or make-to-stock (MTS). Within the MTO environment, there are sub-categories, including Assemble to Order (ATO), Configure to Order (CTO), and Engineer to Order (ETO).

Assemble-to-Order   An environment in which a product or service is assembled after receiving an order from the customer. The key components that are used in the assembly or finishing process, such as bulk, semi-finished, intermediate, subassembly, fabricated, purchased, and packaging, are planned and generally stocked in anticipation of a customer order. Receipt of an order initiates the assembly of a customized product.

Configure-to-Order   An environment in which designs are created by the manufacturer's engineers, but design options are provided for most of the important components that go into the unit. Custom configurations are typically created by including different combinations of standard items.

Engineer to Order   An environment in which products are specifically designed to the specifications of the customer. Sub-Assemblies components can be either standard items or custom items that are designed specifically to the customer order. It is common for ETO environments to utilize mixed mode manufacturing strategies, including elements of MTS, ATO, CTO, and MTO processes.

Make to Stock   An environment in which items can be finished and usually are finished before receiving a customer order. Customer orders are typically filled from existing stocks and production orders are used to replenish those stocks.

Scheduling Work Orders with Automatically Generated Routings

For scheduling purposes, the business requirement of effectively dealing with MTO environments presents various challenges to an enterprise looking to create optimal plans based on changing manufacturing processes. In addition, companies that are not a pure MTO environment also deal with dynamic manufacturing processes. This can be as simple as issuing a different material to a work order that was not part of the standard ERP routing, charging for additional machines, crews, or tools for which the jobs were not setup to run on, or running additional operations on a part. For scheduling purposes, any work order that is imported into Production Scheduling and varies from the standard template from ERP, other than from pre-designated alternates, can be considered a configured work order.

For cMRO implementations, all routings and operations for maintenance work order are auto-generated and are based on the work order that comes from the cMRO system.

Typical Scenarios

This list outlines some typical scenarios that present a configured work order to Production Scheduling:

*   Work orders, which are generated in a *Configurator* program.

    Essentially, work orders are dynamically created and passed to Production Scheduling. No corresponding routing exists in Production Scheduling.

*   Maintenance work orders, which are generated from the Complex Maintenance, Repair, and Overhaul (cMRO) system.

*   Non-standard crews, machines, or tools that are included on the work order bill of resource.

*   A work order that is missing crews, machines, or tools on work order operations.

*   Non-standard items that are included on the work order bill of material.

*   Missing items on the work order bill of material

*   Additional operations that are included on the work order routing.

*   Missing operations on the work order routing.

    For example, Oracle EBS does not pass "Closed" work order operations and send to Production Scheduling.

Production Scheduling has the capability to remove the dependency on having a routing template pre-created in the model for every work order in the application. The application dynamically cross references the work order routing against the routing template during the import process. If Production Scheduling comes across one of the problematic scenarios listed above, it creates a new routing, which directly corresponds to the work order in question. If operations differ from their standard templates, Production Scheduling creates a new operation, which is present in the routing.

Example of a Work Order That Results in a Routing

In this example, assume that the following routing exists in Production Scheduling as a standard template:
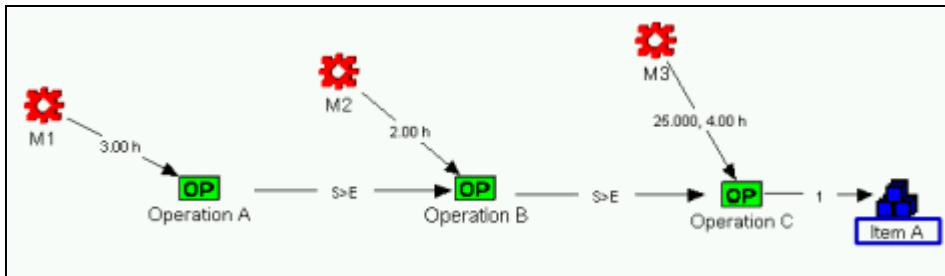
Standard routing template

Now, assume the work order which is created in ERP contains an additional machine "Mill" on Operation C, which is a non-standard machine for this part.
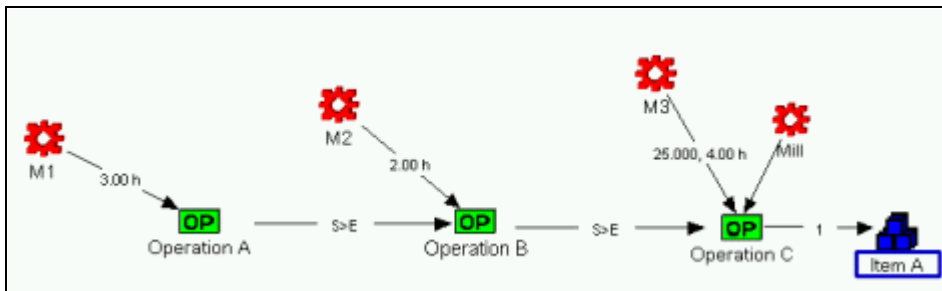
Additional operation on Machine C

When this work order is imported into Production Scheduling, a subfolder of the All Routings folder is created. This folder is titled *Auto-Generated,* as shown in the diagram below. The Auto-Generated subfolder contains the newly created routing.
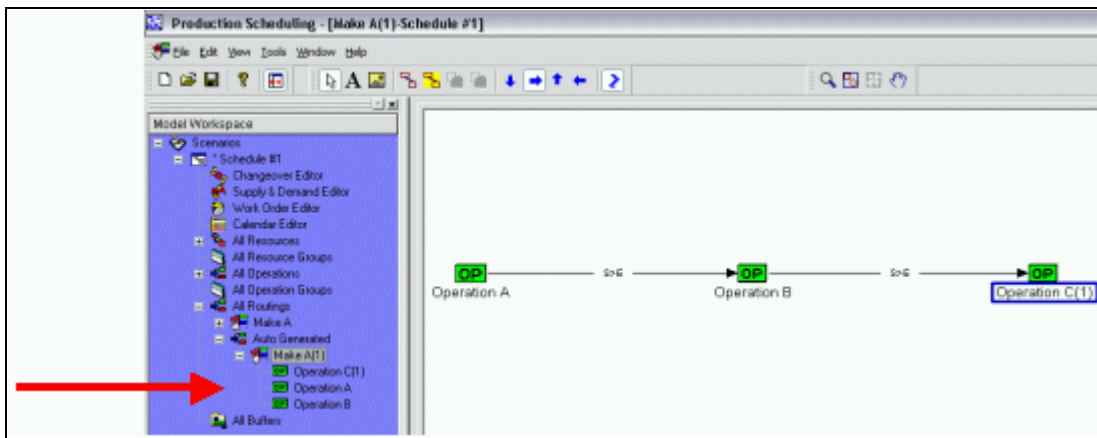
Subfolder created in All Routings

In addition to the new subfolder, the following events occur:

1. A new routing exists in Production Scheduling. The name is specified on the work order header. If the name already exists, Production Scheduling appends the number one in brackets, (1), to the routing code. The integration ID from the XML file is retained and sent back when the schedule is published. In this way, the inbound integration can continue to cross-references the original routing that was returned to fill in any relevant details.

2. The routing attribute *Use for Work Orders Only* is populated. No net new production recommendations use this routing to create new production when total demand is grater than the sum of work order quantities.

3. The sub-folder of the All Operations folder named Auto-Generated, contains a newly created Operation C(1). This is the operation that had the additional resource charged to it. If the name is the same as an existing operation, Production Scheduling appends a number one in brackets, (1) to it's operation code. The integration ID from the XML is retained and sent back when the schedule is published so the inbound integration continues to function, as shown in the diagram below:
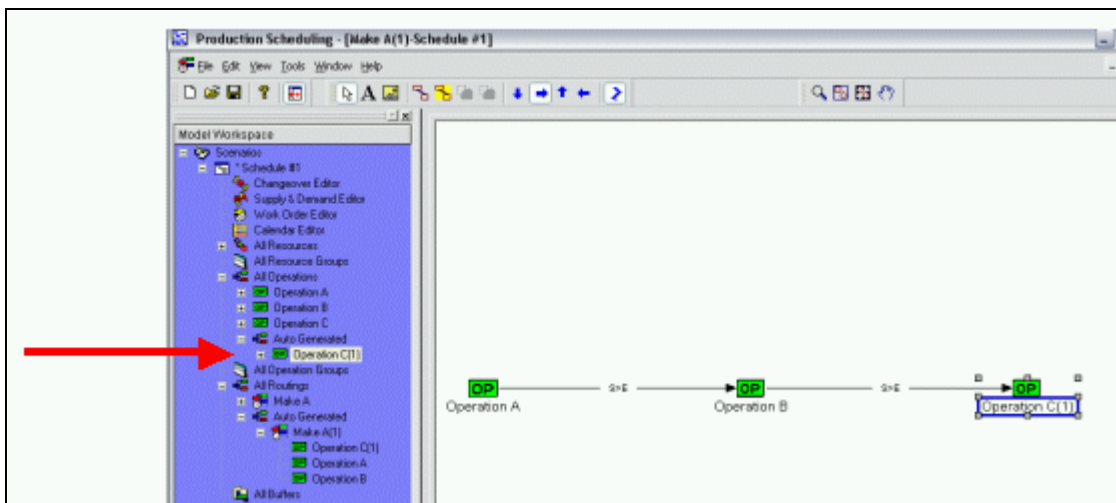
New operation folder appended with (1)

This diagram illustrates that Operation C is redefined:
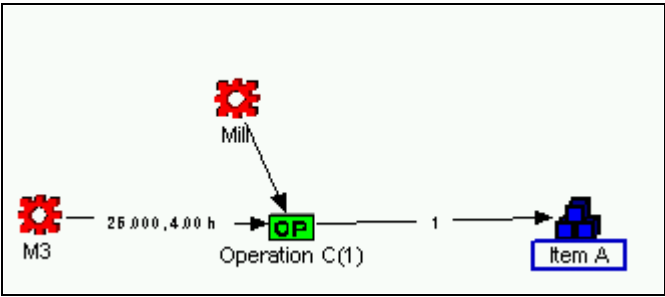
Operation C redefined

This page illustrates that the work order refers to the new routing, contains the new operation code, and additional resources:

New Routing

This page illustrates the additional resources:

Additional resources

This page illustrates the scheduled work order:
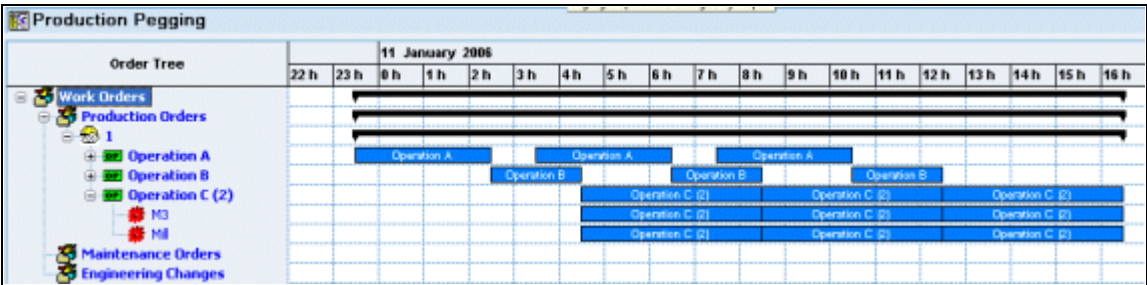
Scheduled work order

Automatically Generated Scenarios and Corresponding Results on Import

The following table outlines several scenarios that result in routings to be auto-generated and, or, operations:

Scenario	Result
A work order is created in a *Configurator* program, which contains no corresponding routing in the Production Scheduling data model.	A routing is created and assigned the code that is present on the work order. All operations that are not previously defined are created in the Production Scheduling data model.
Non-standard crews, machines, or tools are included on the work order bill of resource.	A routing is created in Production Scheduling. Each operation has a new operation created in the Production Scheduling data model. If there is an integrationID assigned to the routing code in the model, it is assigned to this routing code. This is the same for any operation created.

Scenario	Result
The work order operation is missing crews, machines, or tools.	A routing is created in Production Scheduling. Each operation has a new operation created in the Production Scheduling data model. If there is an integrationID assigned to the routing code in the model, the integrationID and the new operations are assigned to this routing code.
Non-standard items included on the work order bill of resource.	A routing is created in Production Scheduling. Each operation has a new operation created in the Production Scheduling data model. If there is an integrationID assigned to the routing code in the model, the integrationID and the new operations are assigned to this routing code.
The work order operation is missing items.	A routing is created in Production Scheduling. Each operation has a new operation created in the Production Scheduling data model. If there is an integrationID assigned to the routing code in the model, the integrationID and the new operations are assigned to this routing code.
Additional operations are included on the work order routing.	A routing is created in Production Scheduling. Each operation has a new operation created in the Production Scheduling data model. If there is an integrationID assigned to the routing code in the model, the integrationID and the new operations are assigned to this routing code.
There are missing operations on the work order routing.	A routing is created in Production Scheduling, which contains only the operations on the work order routing. If there is an integrationID assigned to the routing code in the model, it is assigned to this routing code.

Understanding Work Order Yield and Scrap

Work orders in Production Scheduling represent both yield and scrap on a given operation and across operations. Yield applies to operation output and scrap applies to consumed items.

These examples illustrate yield and scrap on operations.

* Example of yield calculation.

* Example of consumed items.

Example of Yield Calculation

The following criteria is used in this example:

* Work order quantity = Qw.

* Work order routing or operation yield = y.

- Work order parts list scrap factor = s.

- Work order parts list scrap factor quantity = Qc.

This diagram illustrates a manufacturing process, which consists of three operations. Each operation has a quantity of material going into the operation, and a quantity, as a percent yield, coming out.
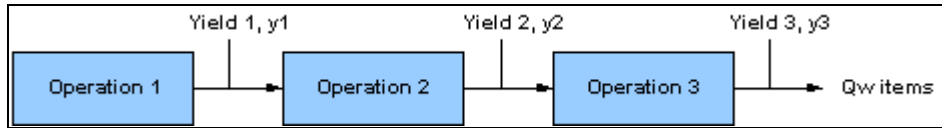
Basic three step manufacturing process with yield

As a result of each operation, there is material that is unusable. This is represented by (1−y) x Qw (quantity of consumed items) and is expressed as a percent. In this example, we want to produce four bicycles. The yield for each operation is as follows:

- Operation 1 yield is 80 percent

- Operation 2 yield is 100 percent

- Operation 3 yield is 50 percent

Operation 1 starts by trying to make 10 bikes. The yield for operation 1 is 80 percent so only 8 bikes are successfully produced. That is, 0.8 x 10=8. Eight good products are successfully produced. Two items are unusable.

Operation 2 continues the process with 8 good bikes because the yield is 100 percent, 1x 8=8.

Operation 3 is the last operation in the process. The yield for this operation is 50 percent so 4 good bikes are produced, 0.5 x 8=4. Four good products are produced. Four items are unusable.

Of the ten bikes that were started in operation 1, the routing delivers four good bikes. The cumulative yield is y1 x y2 x y3 = 0.8 x 1 x 0.5 = 0.4 This is illustrated in the following diagram:
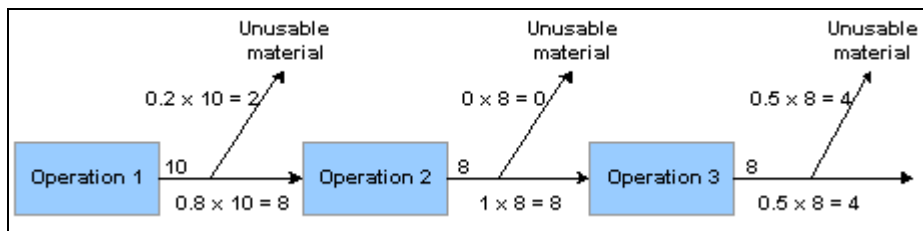
Yield calculation

Example of Consumed Items

To make the bicycles in this example, each bike needs one handlebar, one seat, two wheels, and one drive train. The operations and their scrap factor are:

- Operation 1 consumes one handlebar per bike, scrap is 17 percent

- Operation 2 consumes one seat per bike, scrap is 0 percent

- Operation 3 consumes two items:

  - two wheels per bike, scrap is 20 percent

  - 1 drive train per bike, scrap is 20 percent

Operation 1 starts by trying to consume 12 handlebars. The scrap for operation 1 is 87 percent so only 10 handlebars are successfully consumed. That is, (1−0.17) x 12=10. Two items are unusable, leaving 12 good products.

Operation 2 starts by trying to consume 8 bicycle seats. The scrap is 100 percent, (1−0) x 8, 8 seats are consumed.

Operation 3 starts by trying to consume 20 wheels. The scrap for this operation is 80 percent so 16 good wheels are consumed, (1−0.2) x 20=16. The same scrap factor applies to making drive trains. It starts by consuming 10. Eight are successfully consumed (1 x 0.2) x 10 = 8.

This is illustrated in the following diagram:
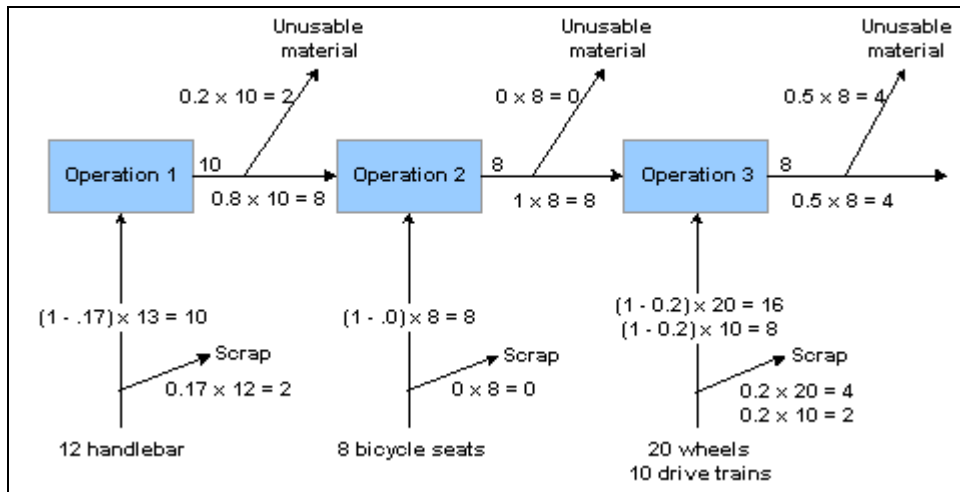
Consumed items calculation

Numerically, the example looks as follows, where:

- AP is Attempted Production

- TP is Targeted Production (4 bicycles)

- QA is Quantity per Assembly

- QRFI is Quantity Removed From Inventory

- S is scrap

- Y is yield

Operation	Consumed Item	Quantity per Attempted Production	Yield (percent)	Scrap	Quantity Removed from Inventory Formula	Quantity Removed from Inventory
3			50		AP=TP/Y3	AP=4/0.5=8
	Wheel	2		20	QRFI=APxQA/(1–S)	QRFI=8x2/0.8=20
	Drive Train	1		20	QRFI=APxQA/(1–S)	QRFI=8x1/0.8=10
2			100		AP=(TP/Y3)/Y2	AP=8/1=8
	Bicycle Seat	1		0	QRFI=APxQA/(1–S)	QRFI=8x1/1=80
1			80		AP=((SKP/Y#)/Y2)/Y1	AP=8/0.8=10
	Handlebar	1		17	QRFI=APxQA/(1–S)	QRFI=10x1/0.83=12

Creating a Work Order

Use this page to create a work order:

Work Order page

Note. Fields that are grayed out are either calculated or predetermined. For example, the Lot Multiple field is grayed out because it is populated from the producing operation of the selected routing.

To add a work order:

1. Access the Work Order Editor.

2. Right-click the folder or the work order category where you want the new work order to reside, and select the appropriate work order type:

   • Production Work Order

   • Maintenance Work Order

   • Engineering Work Order

3. Double-click the new work order.

4. Complete the fields in the main window of the Work Order Editor.

5. Complete the fields under the Routing, Bill of Resources, Bill of Materials, and Produced Item tabs.

6. Click Save and exit from the Work Order Editor.

Work Order Header

The fields and descriptions in the Work Order header are:

Work Order Number	Enter a number for the work order.
Status	Displays the current status of the work order: Valid values are: • *Planned*: the work order is created but not yet released to production. • *Released*: The work order is released to the shop floor for scheduling.
Item Code	Select an item from the list.
Type	Displays the type of work order you have chosen to add.
Order Class	If the demand is pegged, the order class is pre-assigned and assumes the class order of the pegged demand. If demand is not pegged, select one of these options: *Hot,Committed,* or *Uncommitted.*
Priority	Select the priority of the work order. If the work order is linked to another, it inherits the demand priority from the linked demand order line item.
On Hold	Select to indicate that the work order is on hold.
Unit Of Measure	Displays the unit of measure for the item.
Quantity Required	Enter the quantity of the item that is required to fulfill the work order. This field is not available for Maintenance or Engineering work orders.
Quantity Completed	Enter the number of items that have been produced out of the total quantity that is specified in the work order. This field is not available for Maintenance or Engineering work orders.

Quantity Remaining	Displays the number of items that need to be produced to satisfy the work order. This number is calculated by the system. This field is not available for Maintenance or Engineering work orders.
Created	Displays the date that the work order was created in the MRP program.
Work Order Requested Date	Select the date that the work order is to be completed. When linked to a demand order, it inherits the requested date from the Supply & Demand editor.
Off Hold	Select to indicate the off hold date.
	This field is applicable only if the *On Hold* check box is selected. This date becomes the Earliest Start Date on the first operation of the work order. This overrides any firm status on the operations if a conflict exists.
Start	Displays the start date and time of the first operation in the routing. This field is calculated by the Production Scheduling solver.
Completion	Displays the end date and time of the last operation in the routing. This field is calculated by the Production Scheduling solver.
Demand Request Date	Displays the date if a work order is pegged to this demand.
Routing	Select the routing for the work order. This field displays all applicable routings and stand-alone operations that can build the item.
Lot Multiple	Displays the multiplier based on the selected routing.
	Displays the lot multiple of the associated routing. If using units of effort and the work order quantity is greater than the routing lot multiple, the work order is divided into these incremental lot multiples to create several routing instances. For example, if your work order quantity is 4000 units and the routing lot multiple is 2000 units, two routing instances are scheduled.
Original Routing	The Original routing is a reference if the work order is offloaded to an alternate routing by the solver. The option displays the original assignment before offloading occurred.
Demand	Select to link a parent work order to a demand order.
	Note. This field is not available for Maintenance and Engineering work orders.
Line Item	The number assigned to the line item of the demand.
Routing Instances Required	Displays the number of routings needed to satisfy the work order. If Units of Effort is enabled, this displays *1*. If Units of Effort is disabled, this number is the quantity or lot multiple.

Work Order - Routing Tab

The fields and descriptions in the Work Order - Routing tab are:

Operation Code   Displays the operation code for each routing stage. This field comes from the chosen routing.

Operation Description   Displays the description of the operation.

Status   Displays the current status of the operation. Valid values are:

Open: the operation has not yet started.

Active: the operation is in progress.

Closed: the operation is complete.

Next Operation   Displays the name of the operation that follows the current operation. If no operation follows this operation, the field is blank.

Precedence Relationship Select the precedence relationship for this operation from the drop-down menu. This overrides the routing template, if it is different.

See "Understanding Precedence Relationships," Creating Production Scheduling Models, Defining Routings, Working with Routing Diagrams, *Production Scheduling*

Min Separation   Enter the minimum separation time between operations. This overrides the
(minimum separation)   routing template, if it is different.

See "Understanding Precedence Relationships," Creating Production Scheduling Models, Defining Routings, Working with Routing Diagrams, *Production Scheduling*

Max Separation   Enter the maximum separation time between operations. This overrides the
(maximum separation)   routing template, if it is different.

See "Understanding Precedence Relationships," Creating Production Scheduling Models, Defining Routings, Working with Routing Diagrams, *Production Scheduling*

Firm Status   Determines if the start date, end date, or both the start and end date are inflexible. Valid values are:

• *Firm Start*: the start date is fixed. When you select this option, the Start Date field becomes active. Set the start date by moving the up and down arrows or by clicking on the large arrow, which brings up a calendar.

• *Firm End*: the end date is fixed. When you select this option, the End Date field becomes active. Set the end date by moving the up and down arrows or by clicking on the large arrow, which brings up a calendar.

• *Firm Range*: the start date and end date are fixed. When you select this option, both the Start Date field and the End Date field become active. Set the date by moving the up and down arrows in the respective date fields or by clicking on the large arrow, which brings up a calendar.

Start Date   The start date for the operation. This field is populated only if you specify a start or firm range.

End Date	The end date for the operation. This field is populated only if you specify an end or firm range.
Planned Start Date	Displays or Enter the scheduled starting time of an operation as calculated by the solver. When a work order is broken down into its units of effort, this field reflects the start time of the operation instance.
Planned End Date	Displays or Enter the scheduled ending time of an operation as calculated by the solver. When a work order is broken down into its units of effort, this field reflects the end time of the operation instance.
Yield Percentage	Enter the percentage of usable output from the operation. If you change this, Oracle Production Scheduling scales upstream operation durations and bill of material quantities.
Cumulative Yield	The composite yield percentage from the last operation to the first.

Work Order - Bill of Resources Tab

The fields and descriptions in the Work Order - Bill of Resources tab are:

Operation Code	Displays the operation code for each routing stage. This field is populated from the chosen routing.
Planned Resource Set	Displays the preferred resource set.
	You may select a different "All of" set or single resource. When a different All of set is selected, the Planned Resource field reflects the new resource. If the All of Set has a different number of members, rows are added or deleted in the tab to reflect this.
	See "Creating All of Sets," Creating Production Scheduling Models, Defining Operations, *Production Scheduling*
Planned Resource	Displays the resource that was planned to be used to complete this operation.
	Note. This column is display only when the Planned Resource Set field has a value of *All of Set.*
Actual Resource	Displays the resource that was used to complete this operation.
Scale Duration	Select *Yes* to indicate that the machine run time, in the Routing page, is to be scaled proportionally to changes in quantity. For example, if an operation takes 10 minutes to produce 20 units, and the quantity is reduced to 10 units, then proportionally, the reduced time is 5 minutes. Options are: *Yes* or *No.*
	Note. This field is not available for Maintenance and Engineering work orders.
Run Time	Displays the duration of the operation.

Remaining Run Time	Displays the amount of time that is required to complete operations that have a status of Active or Open.

Work Order - Bill of Materials Tab

The fields and descriptions in the Work Order - Bill of Materials tab are:

Operation Code	Displays the operation code for each routing stage. This field is populated from the chosen routing.
Planned Item Code	Displays the item that is planned to be consumed by this operation.
Planned Item Description	A description of the item that is planned to be consumed by this operation.
Actual Item Code	Displays the item that was consumed to complete this operation.
Actual Item Description	Displays the description of the item that was consumed to complete this operation.
Scale Parts	Select whether you want to scale parts. Options are: *Yes* or *No.*
UOM	Displays the unit of measure for the planned item.
Quantity Required	Enter the amount of material required for the operation.
Scrap Percentage	Enter the percentage of input material that is unusable by this operation.
Scrap Adjusted Quantity	Displays the amount of material required, adjusted for scrap.
Quantity Issued	Enter the number of items that remain to be issued or consumed.
Quantity Remaining	Displays the number of items that need to be consumed to satisfy the work order.

Work Order - Produced Items Tab

The fields and descriptions in the Work Order - Produced Items tab are:

Operation Code	Displays the operation code for each routing stage. This field is populated from the chosen routing.
Item Code	Displays the unique identifier of this item.
Type	Displays the type of item, such as a primary item or a by-product.
UOM	Displays the unit of measure for the item code.
Quantity Required	The number of items that need to be produced to satisfy the work order.

Quantity Completed   The number of items that have been produced.

Quantity Remaining   Displays the number of items that still need to be produced to satisfy the work order.

Work Order - Related Work Orders Tab

The buttons, fields, and descriptions in the Work Order - Related Work Orders tab are:

Add (button)   Click this button to select the next work order.

Delete (button)   Click this button to delete the link to a work order that exists in the list or related work orders.

Work Order   Displays the work order number of related work orders.

Type   Displays *Next* or *Previous.*

                              You can only add Next work orders, but the other work orders that are selected list the work order as *Previous.*

Item   Displays the item that is produced from the related work order.

Quantity   Displays the remaining quantity of the requested item from the related work orders.

Request Date   Displays the request date of related work orders.

From Operation   Displays the From Operation of the Previous Work Order.

To Operation   Displays the To Operation of the Next Work Order.

Precedence Relationship   Displays the precedence constraint between the work order header and the related work orders. If you change this field the precedence relationship is updated on the work order header. Valid values are: *Starts After End,Starts At End,* and *Starts After Start.*

Min Separation   Enter the minimum separation time. This field can be changed only if the
(minimum separation)   precedence field has a value of *Starts After End.*

Max Separation   Enter the maximum separation time. This field can be changed only if the
(maximum separation)   precedence field has a value of *Starts After End.*

Completion Date   Displays the completion date of related work orders. This field is populated when the solver runs.

Work Order - CMRO Tab

In environments where Production Scheduling is integrated to cMRO, the fields in the CMRO tab will be populated by default as follows:

Visit	Displays the visit assigned to the work order.
Previous Milestone	The previous milestone that must be completed before this work order can commence.
Next Milestone	The milestone that is reached upon completion of all work orders in the current milestone.
Stage	The stage the work order belongs to
Process	Displays field if cMRO work order uses this field. List displays all the processes assigned to work orders in the model.
System	Displays field if cMRO work order uses this field. List displays all the systems assigned to work orders in the model.
Access Panel	Displays field if cMRO work order uses this field. List displays all the access panels assigned to work orders in the model.
Zone	Displays field if cMRO work order uses this field. List displays all the zones assigned to work orders in the model.
Subzone	Displays field if cMRO work order uses this field. List displays all the subzones assigned to work orders in the model.
Department	Displays field if cMRO work order uses this field. List displays all the departments assigned to work orders in the model.
Unit Name	Displays field if cMRO work order uses this field. List displays all the unit names assigned to work orders in the model.
Non Routine	Displays field if cMRO work order uses this field. Valid options are Yes or No. This option is used indicate whether the work order is of type non routine or not.

Deleting a Work Order

To delete a work order:

1. Access the Work Order Editor.

2. In the Work Order tree, select the work order you want to delete. If you want to delete more than one work order, hold down the Shift key and highlight the appropriate work orders or folders. If the work orders or folders are not listed sequentially, hold down the CTRL key and select the individual work orders or folders.

3. Right-click and select Delete from the menu.

Accessing Views from a Work Order

You can access the Production Pegging, Resource Gantt, or Operation Gantt from the Routing, Bill of Resources, Bill of Materials, or Produced Items tab displayed at the bottom of the Work Order Editor.

To access a view from a work order:

1. Access the Work Order Editor.

2. Expand the necessary folder and double-click a work order.

3. Right-click the Operation Code field from the Routing, Bill of Resources, Bill of Materials, or Produced Items tab and select one of the following: Drill down to Production Pegging; Drill-down to Resource Gantt; or Drill down to Operation Gantt.

   The selected view appears. Performing a drill down from the Work Order Editor highlights the selected operation and performs a context sensitive zoom as well, making it easier to find the work order.

Creating Work Order Folders

To create work order folders:

Note. To expand the Work Order Editor window and show or hide the work order tree, click on the Show/Hide icon in the menu bar.

---

1. Access the Work Order Editor by clicking the icon in the menu bar.

2. Right-click on the type of work order folder you want to create.

   Options are: *Production Orders,Maintenance Orders,* or *Engineering Orders.*

3. Select Add Folder from the drop-down menu.

4. Right-click the folder, select Rename, and enter a new name for the folder.

5. Press Enter.

The Add Folder and Add Production Work Order menu are shown in the graphic below:
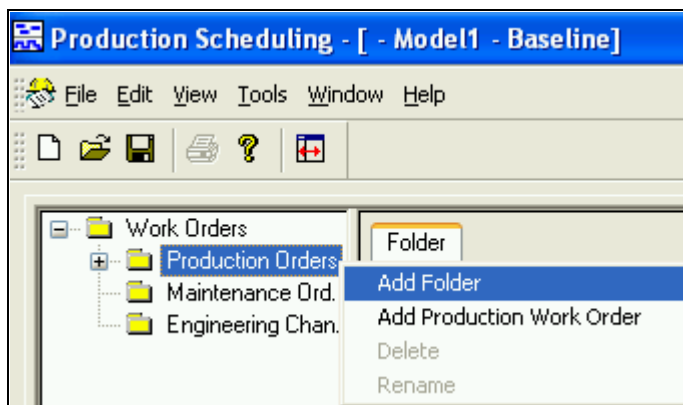
Work Order Editor, Add Folder, and Add Production Work Order menu

Organizing Work Orders Into Folders

To organize work orders into folders:

1. Access the Work Order Editor.

2. Expand the needed folders to see the work orders that need to be moved, and the folders that you wish to move the work orders into.

3. Select the work order or work orders that you wish to move.

4. Drag and drop the work orders into the desired folder.

Creating and Viewing Work Order Folders

To adjust or view work order folder properties:

1. From the Work Order Editor, right-click work order folder from work order tree. Select Change. The Change – Work Order dialog appears.

2. Select the General tab to view or modify work order quantity, request date, and hold status.

3. Select the CMRO tab to view visit information.

4. Click OK to apply and close the dialog box. Click Cancel to close dialog box without saving your changes.

5. If the request date for a milestone is modified, a solve or repair solve is required.

Creating Hard Links Between Work Orders

This section provides an overview of hard links between work orders and discusses how to:

• Create a hard link between work orders.

• View related work orders.

Understanding Hard Links Between Work Orders

The concept of *Next* work orders exists in Production Scheduling. This means that one particular work order must complete before the *next* work can begin. If using the work order units of effort feature, the work orders can overlap but there must be enough inventory produced from the first work order before the start of the first operation instance on the *Next* order that consumes it.

Hard links between work orders enables a user to create a many to many relationship between work orders and provides the ability to specify precedence constraints between these work orders. When the user creates the Next relationship, the details are populated for both work orders in the Related Work Orders tab of the Work Order Editor.

The Related Work Orders tab of the Work Order Editor displays the type of relationship (Next or Previous), the item it produces, quantity, request date, the precedence relationship, the minimum and maximum separation, and a completion date. The Precedence Relationship, Min Separation, and Max Separation fields are open fields. The Completion Date field is populated when the solve completes.

It is possible to have multi-level work orders. For example, a Next work order is also Previous to another work order.

If a user creates a Next relationship between two work orders and the inventory quantity between the two work orders is unbalanced, for example, work order one creates 40 units, but work order two consumes 50 units of the same inventory. Then the solver instantiates net new production to avoid the infeasibility.

Creating a Hard Link Between Work Orders

To create a hard link between work orders:

1. Access the work order tree.

2. Select the work order that is to become the Next work order.

3. Drag the Next work order and drop it onto the Previous work order.

   This action creates a shortcut between the work orders, and the work order tree structure remains intact.

   This action updates the Related Work Orders tab within the Work Order Editor.

Viewing Related Work Orders

Use this page to view related work orders:

Related Work Orders tab

For button and field definitions:

See "Creating a Work Order," Creating Production Scheduling Models, Working With Work Orders, *Production Scheduling*

Releasing Work Orders to Production

This section provides an overview of selective work order release and the Release To Production window, and discusses how to:

• View the Release To Production window.

• Release work orders to production.

Understanding Selective Work Order Release

After a schedule is created, optimized, and approved, the master scheduler might choose to pass all the existing work orders to the shop floor and ERP system or decide to selectively release a subset of work orders.

Understanding the Release To Production Window

The Release To Production window provides read-only detailed information about all the planned orders and work orders in the model.

The publish horizon of the selected publish profile is displayed as the Release Horizon at the top of the view. All work orders that fall within the publish horizon are selected when you open the Release To Production window. All work orders in the model with start dates within the release time frame are present in the window and are automatically selected.

The Select All or Deselect All buttons enable the scheduler to make mass changes prior to publishing the schedule. The columns are resizable by dragging one of the borders of the column and users can sort the information in the table based on any chosen column by clicking on the column header. When a user chooses to sort based on a field that contains non-unique values a secondary sort is performed based on the Start Date value. A small arrow on the right side of the header field indicates the current sort criteria.

If you extend the publish horizon after selecting the work orders in the view, the original selection is cleared and you have to make your selection again. The work orders that have start dates within the new release horizon are automatically selected. If you run a solve, all selections are reset.

Viewing the Release to Production Window

Use this page to view the Release to Production window:

Release To Production window

To view the Release to Production Window:

1. Approve the scenario.

2. Right–click on the scenario and select Edit Work Order Release.

   Note. If the scenario is not approved, the Edit Work Order Release option is not available.

3. To view only those work orders that meet a defined filter, select the appropriate filter from the Work order filter list. The window refreshes to display only those order that meet the filter criteria.

Releasing Work Orders to Production

To release work orders to production:

1. Access the Release To Production window.

2. Select one or more work orders, that are ready to be released.

   • Select an individual work order by selecting the box to the left of the row that displays the work order information.

   • Highlight multiple work orders and click the Select All button.

3. Click the OK button.

4. Publish the profile.

    See "Publishing Schedules," Publishing Data, *Production Scheduling*

Defining Changeover Rules

This section provides an overview of changeover rules and discusses how to:

*   Create changeover rules.

*   Re-order changeover rules.

*   Find matching changeover rules.

*   Filter changeover rules.

Understanding Changeover Rules

You can record precise information about sequence dependent changeover time by using the Changeover editor. For example, when a machine die is changed to a larger bore or when a printing press changes to a different color, the machine has to be stopped and reconfigured. Using this feature, you can create attribute based rules that record those times when a machine cannot operate because it has to be set up for the next operation. Costs can also be recorded for each changeover rule. These costs are used to calculate the cost-based metrics in the Key Performance Indicators window.

A hierarchical order exists for the rules that are entered in the Changeover editor. When solving the schedule, the system uses the first rule on the list that applies to the changeover from one machine to another. For example, if you have two rules for the Chain Machines Group from Gold to Bronze, the schedule uses the top rule.

With attribute-based changeovers, it is possible to have one or more rules applicable for a changeover. For example, in a soft drink bottling plant changing the flavor and changing bottling size might have different changeover rules. Because of this it is required to specify whether the changeovers are concurrent or sequential. If they are concurrent, then both changeovers can occur at the same time and the maximum time and cost are used for the changeover rule. Sequential changeovers rules are additive in terms of cost and time. Different calendars can also be applicable for multiple changeovers rules so it is possible to define whether all calendars must be up or just one calendar must be up for scheduling changeovers.

Note. If you create a rule at, or move a rule to the top of the Changeover editor, this rule effectively overwrites all of the other rules.

To save time, you can select groups of attributes or machines when entering changeover rules. For any group other than All Machines or All Attributes to appear, its property attributes must be set to display in the Changeover editor. You can use the filters at the bottom of the view to find out which changeover rule is applicable to the resource or operation to and from combinations.

Creating Changeover Rules

To create a changeover rule:

1. Access the Changeover editor.

2. Click Add new rule.

3. Complete these fields:

Machine	Select the machine, group, or wildcard (*) for which a changeover entry is being made.
Attribute Type	Select Operation or Group, an attribute, or wildcard (*) for which a changeover entry is being made.
From	Select an operation, attribute value, group, or wildcard (*).
To	Select an operation, attribute value, group, or wildcard (*).
Duration	Specify the amount of time that it takes to set up the machine between the From and the To operation changeover. Select the time value (0 hours by default) and specify the expected time for the changeover. To enter other time values, use d for days, h for hours, m for minutes, or s for seconds. You must select, and then change the digit and the letter separately.
Calendar	Select a calendar that applies to the changeover from the list.
Cost	Specify the cost associated with the changeover. When the changeover rule is applied, this cost value is used to calculate KPIs. The Campaign Run Optimization algorithm also uses this cost when determining an ideal sequence dependant setup.

Re-ordering Changeover Rules

To re-order changeover rules:

1. Access the Changeover Editor.

2. Click the row number of the rule that you want to move.

3. Click Move rule up and Move rule down button from the Changeover toolbar to move the changeover rule to its new location.

Finding Matching Changeover Rules

You can use the Find Matching Changeover Rule region at the bottom of the Changeover Editor window to locate changeover rules that meet your criteria.

In the Find Matching Changeover Rule region, perform the following:

1. Access the Changeover editor.

2. Use the Machine list to find changeover rules based on a specific machine.

3. Use the From Operation list to find changeover rules based on a specific From Operation attribute.

4. Use the To Operation list to find changeover rules based on a specific To Operation attribute.

    The Found field displays the number of Changeover rules that match your search criteria. The Duration and Cost fields display the total duration and cost for all changeover rules that meet your search criteria if multiple changeovers are defined as sequential in the Attributes tab of the resource Properties page. The maximum cost and duration appear if multiple changeovers are defined as concurrent.

Filtering Changeover Rule

You can use the filter feature to locate changeover rules. Perform the procedure below to use the filter feature.

1. Access the Changeover editor.

2. Click the Toggle Rule Filter button. The Filter row appears at the top of the Changeover Editor and the Toggle Rule Filter button appears highlighted to indicate the filter option is enabled.

3. Click in each cell within the Filter row you that you want to define, and select the appropriate menu option. You can filter on the following fields: Machine, Attribute Type, To, and From.

    Select (Show All) to display all options available for the filter field. Select the * option to display changeover exceptions you have defined using the * feature.

    When a filter cell is selected, a list appears and displays all criteria for that field. As each filter field is assigned, the Changeover Editor window is updated to display the changeover rules that meet the criteria defined by the filter fields.

4. To turn off the filter option, click the Toggle Rule Filter button. The Filter row at the top of the Changeover Editor is no longer displayed.

Chapter 6

Setting Up Global Options

This chapter discusses how to:

*   Set up system logging.

*   Set up diagrams.

*   Specify supporting documents.

*   Set up global settings.

Setting Up System Logging

This section discusses how to configure log files to capture scheduling information, performance statistics, and any validation warnings that may occur during a schedule solve.

Configuring Log Files

Access the Options window. Select the Logging tab.

To configure log files:

Messages	Select this option if you want to log system messages. If you do not select this option, no logging occurs.
Performance Statistics	Select this option to capture application performance statistics in the log file. Timing for loading, saving, and duplicating schedules are captured.
Show validation warnings	Select this option to have the system to notify you of noncritical issues with data being sent to the Production Scheduling solver. If you select this option, this information is also written to the log file.
Log file	Specify a location to store the log file.
Overwrite log file	Select this option to have the log file overwritten each time that the system runs. To keep old log records, save the log file with a different name each time that you start the system.
Maximum file size (KB)	Specify a file size or accept the default value of 1000 KB.

Truncate file at (%)   Specify the percentage of the log file that the system omits when the log file
                                  reaches the size in the Maximum file size field.

Lowest info display level Select a value from 0 to 9 to specify the level of detail about system performance
                               that is included in the log file. A value of 0 indicates the lowest level of detail,
                               and 9 indicates the highest level of detail.

Setting Up Diagrams

This section discusses how to configure the appearance of operation and routing diagrams to suit your
individual requirements.

Configuring Operation and Routing Diagram Layout

Access the Options window. Select the Diagram Layout tab.

To configure diagram layout:

Top to bottom   Select this option to display operations and routings vertically beginning at the
                                 top of the screen.

Left to right   Select this option to display operations and routings horizontally beginning at the
                                 left of the screen.

Bottom to top   Select this option to display operations vertically beginning at the bottom of the
                                 screen.

Right to left   Select this option to display operations horizontally beginning at the right of the
                                 screen.

Use snake layout   Select this option to display a routing diagram that reads from left to right, and
                                 then from right to left on the next row. This option makes long routings easier to
                                 view on the screen or in printed form.

Specifying Supporting Documents

This section discusses how to create a document directory to store supporting documents for various elements
in your supply chain.

Creating a Document Directory

Access the Options window. Select the Documents tab.

To specify supporting documents:

Documents base folder   Specify the default directory for documents, or browse your system.

Setting Up Global Settings

This section discusses how to set up global settings.

Setting Up Global Settings

To set up global setting:

1. Access the Options window.

2. Select the Global Settings tab.

3. Complete these fields:

Maximum history states   Enter a number that sets the maximum number of history states retained when using the Undo/Redo feature.

Number of rows in bubble help   Enter a number that sets the maximum number of rows for the Notes field in the schedule editor bubble help boxes. Notes are not included in the bubble help when this value is zero.

Apply Oracle Look and Feel   Select to apply an Oracle look and feel to the Production Scheduling user interface. Additional changes to the system desktop colors are required for maximum effect. This procedure is documented in the README file supplied with the Production Scheduling installation.

Allow Drag and Drop in Gantt Views Select to indicate that manual scheduling, such as cut and paste operations, in the Schedule Editor may be performed by dragging selected scheduled operations to their new location. Valid paste locations are highlighted once the dragging operation are started.

Enable Supply Tolerance Offset Select to indicate that supply tolerance offset should be used. If this option is selected, there will be no net new production recommendations generated at any level in the BOM (for example, only work order supply production will be scheduled). If there is insufficient work order supply, inventory will be automatically added to prevent levels from falling below zero (demand is permitted to be shorted). A Production Scheduling log message is generated to indicate what inventory quantity adjustments were made. Using this feature assumes work order supply is sufficient to meet demand.

Chapter 7

Configuring Scenario Properties

This chapter discusses how to:

- Access the Scenario Properties window.

- Set up the Scenario Properties - General page.

- Set up the Scenario Properties - Horizon page.

- Set up the Scenario Properties - Unit of Measure page.

Accessing the Scenario Properties Window

To access the Scenario Properties window:

1. From the Model Workspace, right-click on a scenario.

2. Select the Properties option.

Setting Up the Scenario Properties - General Page

Use this page to set up general model information:

Scenario Properties - General page

To set up scenario properties:

1. Access the Scenario Properties window.

2. Complete these fields:

Code	Specify a name for the schedule.
Branch/Organization	Enter the branch code or organization code for your schedule.
KPI Cost Parameters (key performance indicator cost parameters)	Select a currency symbol. If the desired symbol does not appear, you can enter it in the space provided.
	In the Cost Unit field, select one of the following base units to use when measuring costs: *Minute, Hour, Shift, Day,* or *Week.* This field provides the cost time unit for resource and operating costs.
Use Theory of Constraints Principles	Select this option to enable Theory of Constraints functionality.
	See "Working with Buffers," Creating Production Scheduling Models, Defining Routings, *Production Scheduling*
Notes	Specify relevant information about this model or schedule. For example, if you have made what-if model changes, you can describe these changes here.

Setting Up the Scenario Properties - Horizon Page

The Scenario Properties - Horizon page is use to set up the schedule horizon.

To set up the Scenario Properties - Horizon page:

1. Access the Scenario Properties window.

2. Select the Horizon tab.

3. Complete these fields:

Start Date	Specify the start date of the schedule horizon.
End Date	Specify the end date of the schedule horizon.
First Cycle Start Time	Specify the start date for the first production cycle. Repetitive manufacturing optimized resources use this time to identify the beginning of a production cycle.

Setting Up the Scenario Properties - Unit of Measure Page

Use this page to set up valid units of measure for the scenario:
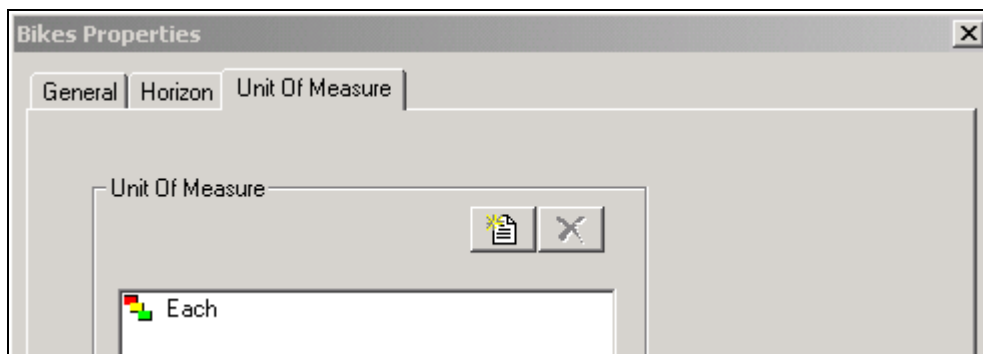
Scenario Properties - UOM page

To set up the Scenario Properties - Unit Of measure page:

1. Access the Scenario Properties window.

2. Select the Unit of Measure tab.

3. Click the Add button to add a unit of measure.

After a unit of measure is defined you can:

• Click the Delete button to delete a unit of measure.

• Right-click and select Delete to delete a unit of measure.

• Right-click and select Rename to rename a unit of measure.

• Associate the unit of measure to an item and the applicable units of measure appear in applicable views.

Note. The *each* unit of measure cannot be deleted, as you need to have at least one in the model.

All data for an item must be in the same unit of measure. For example, you cannot manufacture an item in each's and order the item in cases. Production Scheduling does not perform unit of measure conversion.

Chapter 8

Configuring Solver Options

This chapter discusses how to:

* Access scenario specific solver options.

* Set up general solver options.

* Set up solve sequences.

* Set up campaign run optimization.

Accessing Scenario Specific Solver Options

To access the Solver Options page:

1. From the Model Workspace, right-click on a scenario.

2. Select Solver Options.

Setting Up General Solver Options

This section provides an overview of general options, scheduling work orders according to units of effort, resource criticality, and discusses how to set up general solver options.

Understanding General Options

Production Scheduling can govern the scheduling of operations that exist upstream from a bottleneck operation in a routing. For example, consider a routing with five operations: saw, weld, heat, assembly, and test and pack. The bottleneck resource is the oven, which is required for the heating operation. The routing and operations are displayed in the following diagram:
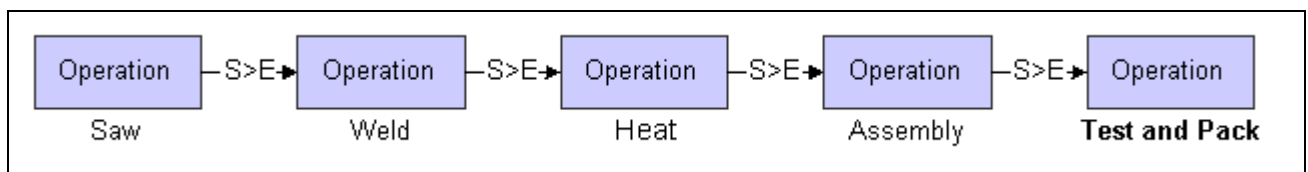
Just in time (JIT) routing and operations

When Production Scheduling creates operations to satisfy demand, it posts earliest date and available date constraints that are associated with the operations. The operations are scheduled in a just in time (JIT) manner, according to their available date. This scheduling behavior is governed by any precedence constraints in the routing and can occur only if precedence constraints permit JIT to available date.

In this example, the saw and weld operations could have been made earlier in the process, however, these operations are scheduled to align with the bottleneck operation, heat. This is the default behavior in Production Scheduling.

Due to the bottleneck created by the oven, the demand is scheduled late. A solve yields the schedule displayed in the following diagram:
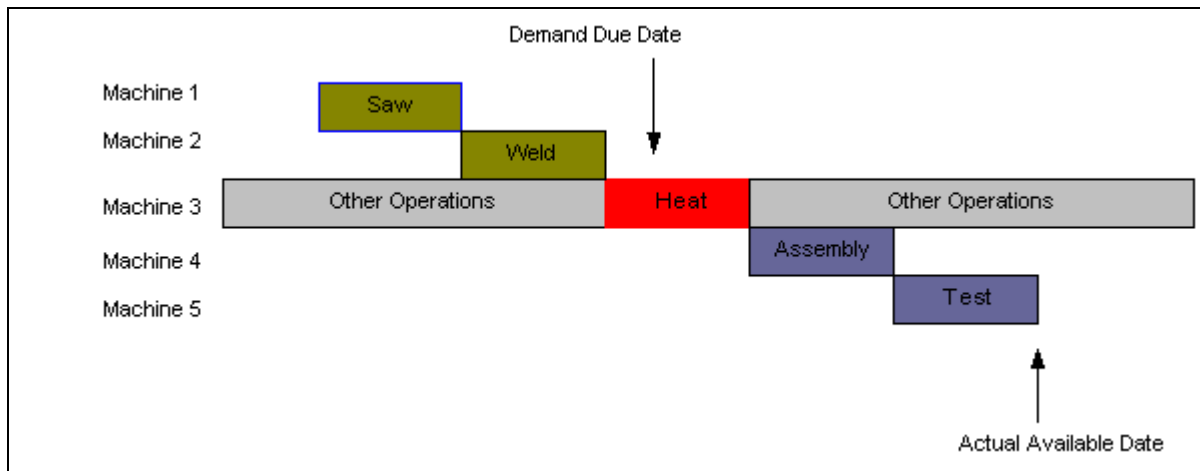
Schedule for demand is late

Using the solver option, JIT to Due Date, you can pull the saw and weld operations forward and generate a schedule resembling the one in this diagram:
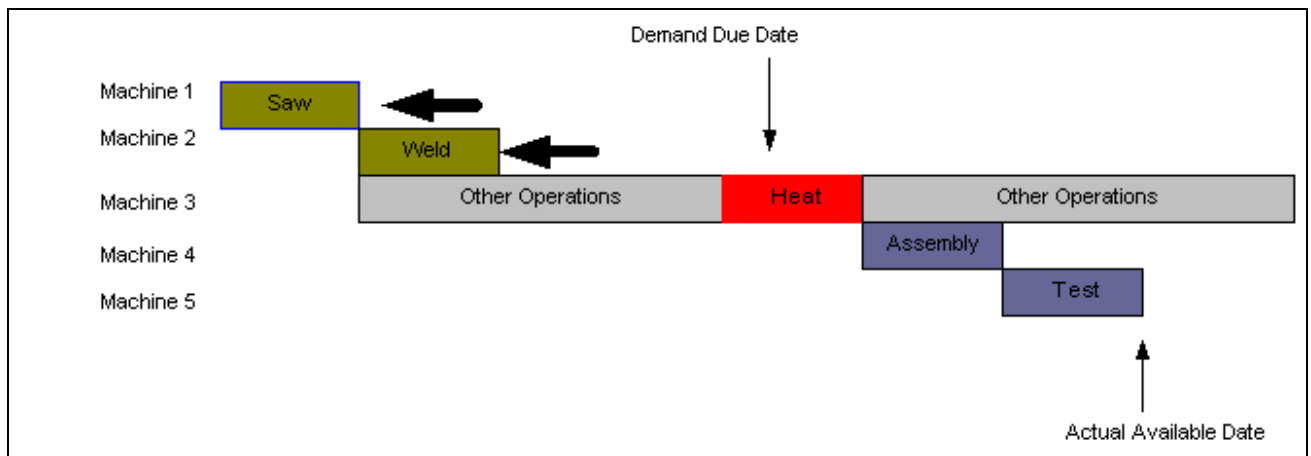
Upstream operation pulled forward

With the saw and weld operations pulled forward, the upstream operations are placed at a point that provides you with a more efficient way to manually schedule the remaining operations in the routing to a time earlier in the process. This helps to make the demand on time and is effective if you want to schedule operations close to the demand due date rather than close to the bottleneck operation on the routing.

Note. This tactic may result in a longer makespan for a routing if this solver option is set to Due Date. For example, in the example, if the solver places upstream operations earlier in the scheduling horizon and if the solver cannot move the bottleneck operation without manual intervention, the makespan is longer for this demand.

Understanding Scheduling Work Orders According to Units of Effort

This section provides an overview of scheduling work order operations and discusses these scenarios:

1. A schedule that has the Units of Effort option disabled on the routing.

2. A schedule that has the Units of Effort option enabled and adjacent operations disabled on the routing.

3. A schedule that has both the Units of Effort and the adjacent Operations Preferred options enabled on the routing.

Scheduling Work Order Operations

In each of these examples, producing Item A requires three operations: Operation A, Operation B, and Operation C. Operation A runs on Machine 1 for three hours; Operation B runs on Machine 2 for one hour; Operation C runs on Machine 3 for two hours. The lot multiple is 20 and Item A takes six hours to produce in this schedule. The routing is shown in the following diagram:
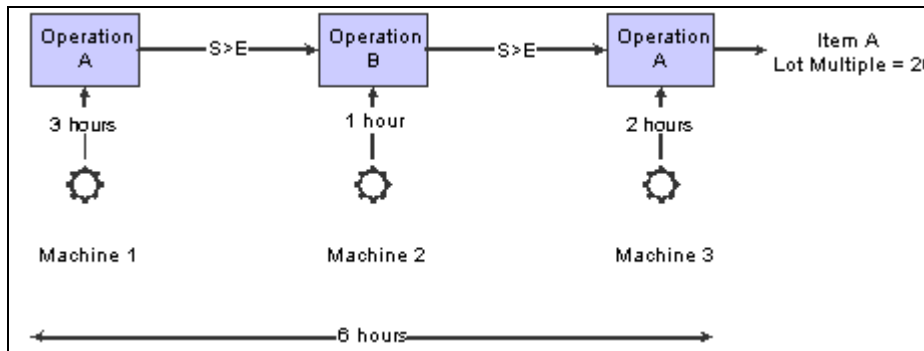
Basic routing for producing Item A

Scenario 1: A Schedule That Has the Units of Effort Flag Disabled on the Routing
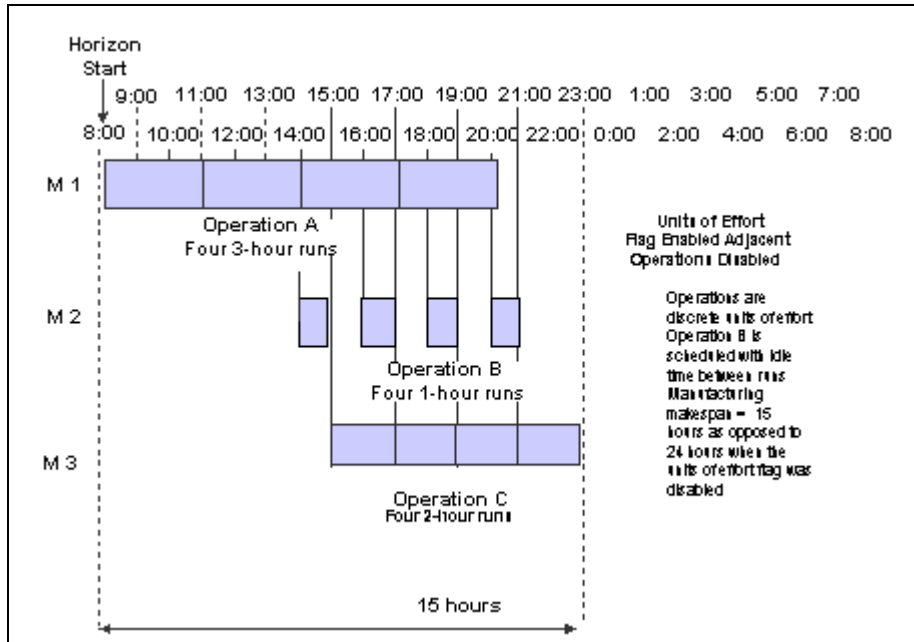
For this scenario, assume that the:

- Order due date: January 1, 4:00 p.m.

- Horizon start: January 1, 8:00 a.m.

- Order quantity: 80

- Enable Units of Effort: NO

The number of hours that each machine can run per operation is:

Machine Number	Runs per Operation
Machine 1	1–12 hour operation
Machine 2	1–4 hour operation
Machine 3	1–8 hour operation

This diagram demonstrates the routing schedule with the Schedule Work Order operation according to the Units of Effort option disabled:
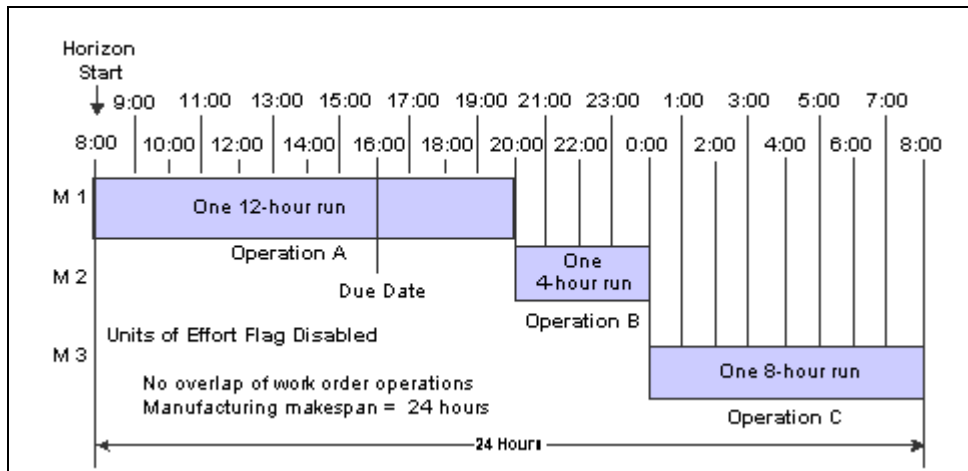
Routing schedule with Units of Effort flag disabled

With this option disabled, there is no overlap on work order operations. It takes 24 hours to produce Item A.

Scenario 2: A Schedule That Has the Units of Effort Flag Enabled and Adjacent Operations Disabled on the Routing

For this scenario, assume that the:

- Order due date: January 1, 4:00 p.m.

- Horizon start: January 1, 8:00 a.m.

- Order quantity: 80

- Enable Units of Effort: YES

The number of hours that each machine can run per operation is:

Machine Number	Runs per Operation
Machine 1	4–3 hour operations
Machine 2	4–1 hour operations

Machine Number	Runs per Operation
Machine 3	4–2 hour operations

The following diagram demonstrates the routing schedule with the Schedule Work Order operation according to their Units of Effort option enabled:

Routing schedule with units of effort flag enabled adjacent operations flag disabled

By enabling the scheduling to apply work order units of effort, substantial benefits are realized:

- Manufacturing makespan is substantially reduced; manufacturing time is reduced from 24 hours to 15 hours.

- Material consumption and production is more accurate.

- Resource utilization is more accurate.

Scenario 3: A Schedule That Has Both the Units of Effort and the Adjacent Operations Preferred Flags Enabled on the Routing

For this scenario, assume that the:

- Order due date: January 1, 4:00 p.m.

- Horizon start: January 1, 8:00 a.m.

- Order quantity: 80

- Enable Units of Effort: YES

- Adjacent operations: YES

Note. You cannot enable adjacent operations unless the Units of Effort option is enabled.

The number of hours that each machine can run per operation is:

Machine Number	*Runs per Operation*
Machine 1	4–3 hour operations
Machine 2	4–1 hour operations
Machine 3	4–2 hour operations

The following diagram demonstrates the routing schedule with both the Schedule Work Order operation according to their Units of Effort and Schedule Adjacent Operations options enabled:
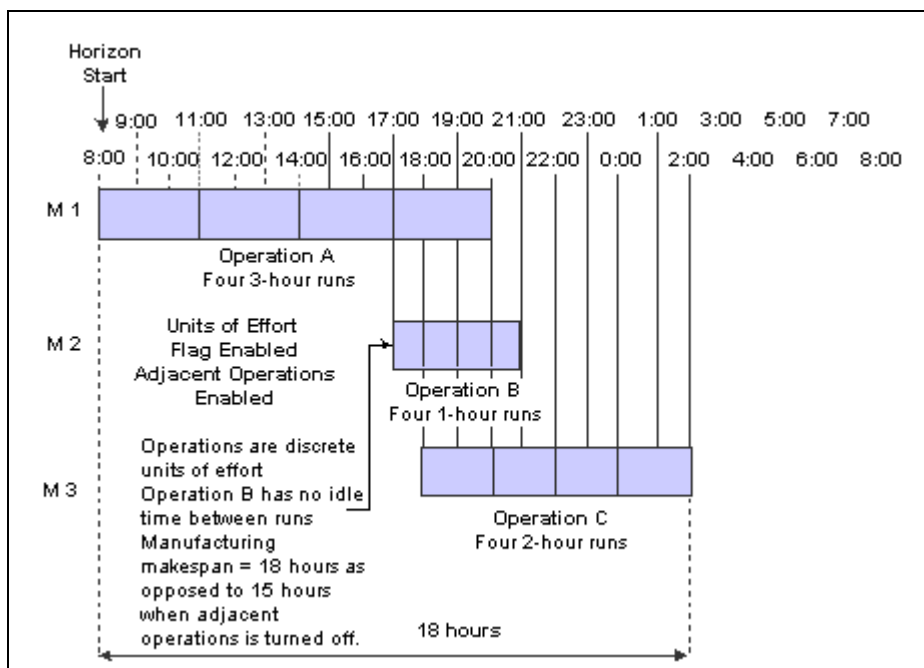
Routing schedule with both Units of Effort flag and Adjacent Operations Preferred flags enabled

By enabling both work order units of work and adjacent operations flags, you have sacrificed some makespan to keep operations on Machine 2 together. Without adjacent operations enabled, total manufacturing time was 15 hours. With adjacent operations on, total manufacturing time is 18 hours, but you have removed idle time on Machine 2.

Planned Start and Planned End Times

Because of the nature of work orders, a given operation within a work order cannot be split into several instances. This can pose problems when operations are broken down into their individual units of work. For example, Operation B becomes Operation B1, Operation B2 and Operation B3. Planned start time and Planned end time fields in the routing tab of the Work Order Editor address this issue. The difference between the two fields is in the logic of how the dates are rolled up into these fields.

When a work order is not broken down into its units of effort, the start and end dates reflect the start and end times of the operation, respectively. However, when a work order is broken down into its units of effort, the planned start date reflects the start time of the first operation instance, Operation B1, and the Planned End Date is the end time of the last operation instance, Operation B3, as illustrated in the following graphic:
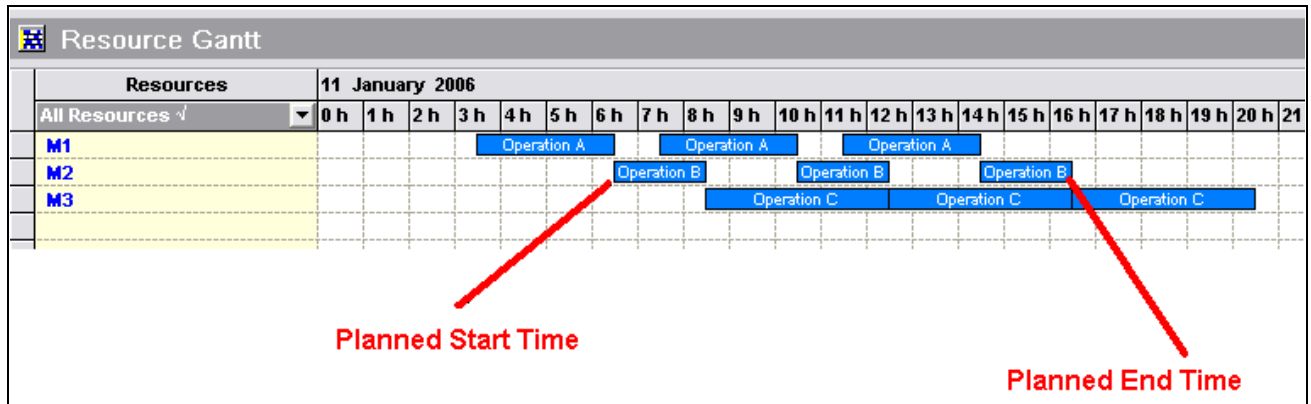
Planned start time and planned end time as reflected in the Resource Gantt

The corresponding work order also reflects those dates as shown in the following diagram:

Planned start time and planned end time as reflected in the corresponding work order

When the Work Order units of effort is enabled, the number of routing instances that are scheduled to fulfill the work order quantity, is displayed in the work order header. The number takes into account the routing lot multiple that is required to perform this calculation. For example, if a work order for a routing is for 50 units and the lot multiple is 25, the quantity is displayed in the work order header.

Note. The number of routing instances that are schedule to fulfill the work order quantity is displayed only if units of effort is enabled. If this feature is not enabled, the field shows a one ("1").

This field always appears as an integer value and is rounded up if a work order is entered in a non-lot multiple quantity.

Understanding Resource Criticality

Production Scheduling has the ability to configure whether the solver first addresses earliest or most resource criticality during solve sequencing. Resource contention, or criticality, is a measurement that occurs in the solver when two or more operations compete for the same resource at a given point in time in the scheduling horizon. The more operations that need to be scheduled on that resource at that given point in time coupled with its capacity, the higher the resource contention or probability that constraints will be violated. In Production Scheduling, you can set whether the solver first addresses earliest contention or most contention resource criticality during solver sequencing.

Focusing and resolving the most critical constraints in the horizon has generally been regarded as the most effective scheduling strategy to resolve floating bottlenecks when sequencing operations. Addressing the more critical problems first results into the problems being broken down into smaller subproblems which are easier to resolve.

Solving in this fashion provides high quality results for most Production Scheduling client models. However, sometimes resolving the highest levels of criticality in the horizon first may not necessarily be optimal from a scheduling perspective, but rather addressing the earliest instances of resource criticality may be more efficient in some cases, such as (but not limited to):

- Routings that contain recursive resources that are considered manufacturing bottlenecks.

- Coupled with the recursiveness, a mixture of crewed and noncrewed operations that also require these bottleneck resources.

- Preferred run times for noncrewed operations during non-business hours (for example, evenings and weekends).

It is recommended to implement this feature following the guidelines above and to run some trials with this option enabled and disabled to understand the differences in schedule quality.

Setting Up General Solver Options

To modify the Solver Options for a scenario:

1. Access the Solver Options page.

2. From the General tab, complete these fields:

Build Strategy	When you define demand policies, you can use the Build Strategy field to configure how Production Scheduling fulfills demand.
	Select *JIT* to indicate that items are scheduled for production in a just in time fashion. When this option is selected, demands are fulfilled with production occurring as close to the due date as possible. The JIT build strategy is useful if you want to avoid storing inventory.
	Select *Pre-Build* to indicate that the solver schedules production at the earliest available opportunity. When this option is selected production is generally started at the beginning of the schedule horizon.
Relax Purchased Items Beyond Lead Time	When this option is enabled, the order multiple is internally set to a very large value and effectively ignored beyond lead time – one will see the addition of a new red line in the item graph to represent the time of the item. Beyond the lead time, the item graph is shaded to indicate infinite supply.
	This option allows you to control whether or not purchase item multiples are respected beyond supplier lead time or not. Keeping these "relaxed" helps solve performance and is generally recommended.
	Relax Purchased Items Beyond Lead Time is set to Yes by default. If the option is set to No, the item graph is displayed taking into account the order multiples.
Scheduling Work Order operations according to their unit of effort	From the drop-down menu, select one of these options: • *Yes* to enable units of effort. This option indicates that Production Scheduling should cross reference the work order to the corresponding routing lot multiple. This results in the work order being divided into lot sized quantities for more flexible sequencing, reduced manufacturing makespan, and accurate item consumption profiles. • *No* to disable units of effort. • *Select by Routing* to schedule selected routings by units of effort. When this option is selected, you must configure the applicable routing properties. See "Setting Routing Properties," Creating Production Scheduling Models, Defining Routings, *Production Scheduling* See "Understanding Scheduling Work Orders According to Units of Effort," Configuring Solver Options, Setting Up General Solver Options, *Production Scheduling*
Adjacent operations referred	Select this option if you want to enable adjacent operations. See "Understanding Scheduling Work Orders According to Units of Effort," Configuring Solver Options, Setting Up General Solver Options, *Production Scheduling*

3. To open the Solve Sequence page, click Solve Sequence.

4. Click OK to save and close the Solver Options page.

Setting Up Resource Business Objectives

This section provides an overview of scheduling resources and groups of resources and discusses how to:

- Access the solve sequence page.

- Create solve sequence stages.

- Change solve sequence stages.

- Remove solve sequence stages.

- Rename solve sequence stages.

Understanding Scheduling Resources and Groups of Resources

From a resource perspective, Production Scheduling generically considers all crews, machines, tool capacity and calendar constraints simultaneously when sequencing operations. Focusing in on and resolving the most critical constraints in the scheduling horizon has generally been regarded as the most effective scheduling strategy to resolve multistage floating bottleneck problems when sequencing operations. Addressing the more critical problems first results in the sequencing problems being broken down into smaller subproblems, which are usually easier to resolve.

It is sometimes desirable to be able to dictate which resources the solver should focus on first when creating a schedule as opposed to having the solver simultaneously consider all constraints and potentially jump between resources when contention levels stipulate this behavior. Supporting resources may be considered secondary in terms of the priority of having their constraints considered during the solve process.

Production Scheduling explicitly allows users to dictate to the solver which resource or group of resources should be focused on and in what order they should be considered. This approach enables scheduling in an incremental fashion by means of some simple configurations in the application. It is also possible to stop the solver after a particular solve stage to review the schedule or make manual scheduling changes.

Determining the priorities by which the solver will address resources or groups of resources is achieved by setting up solver stages. The stages are "lined up" in the order in which the solver will address them.

Resources within a schedule can be assigned to solver sequence stages. Solver sequence stages determine the order in which the solver activates constraints during a solve. Different solve parameters can be specified by stage so, for example, Stage 1 can be solved using the build strategy Pre-Build and the other stages can be solved using JIT. By including resources in a stage, you can set the priorities in which the resources are addressed and solved by the solver. The solver addresses Stage 1, followed by Stage 2, Stage 3, and so on. Resources that are not included in a solve stage are scheduled last.

Finally, with solve stages you can enable multi-stage Campaign Run Optimization.

If no solver sequence stages are set up, the solver considers all resource constraints simultaneously.

Example of Using Solve Sequence Stages

The following example demonstrates scheduling resources and resource groups, using a general flow of product through a factory, starting with the Mounter and ending with the Tape and Reel resources.
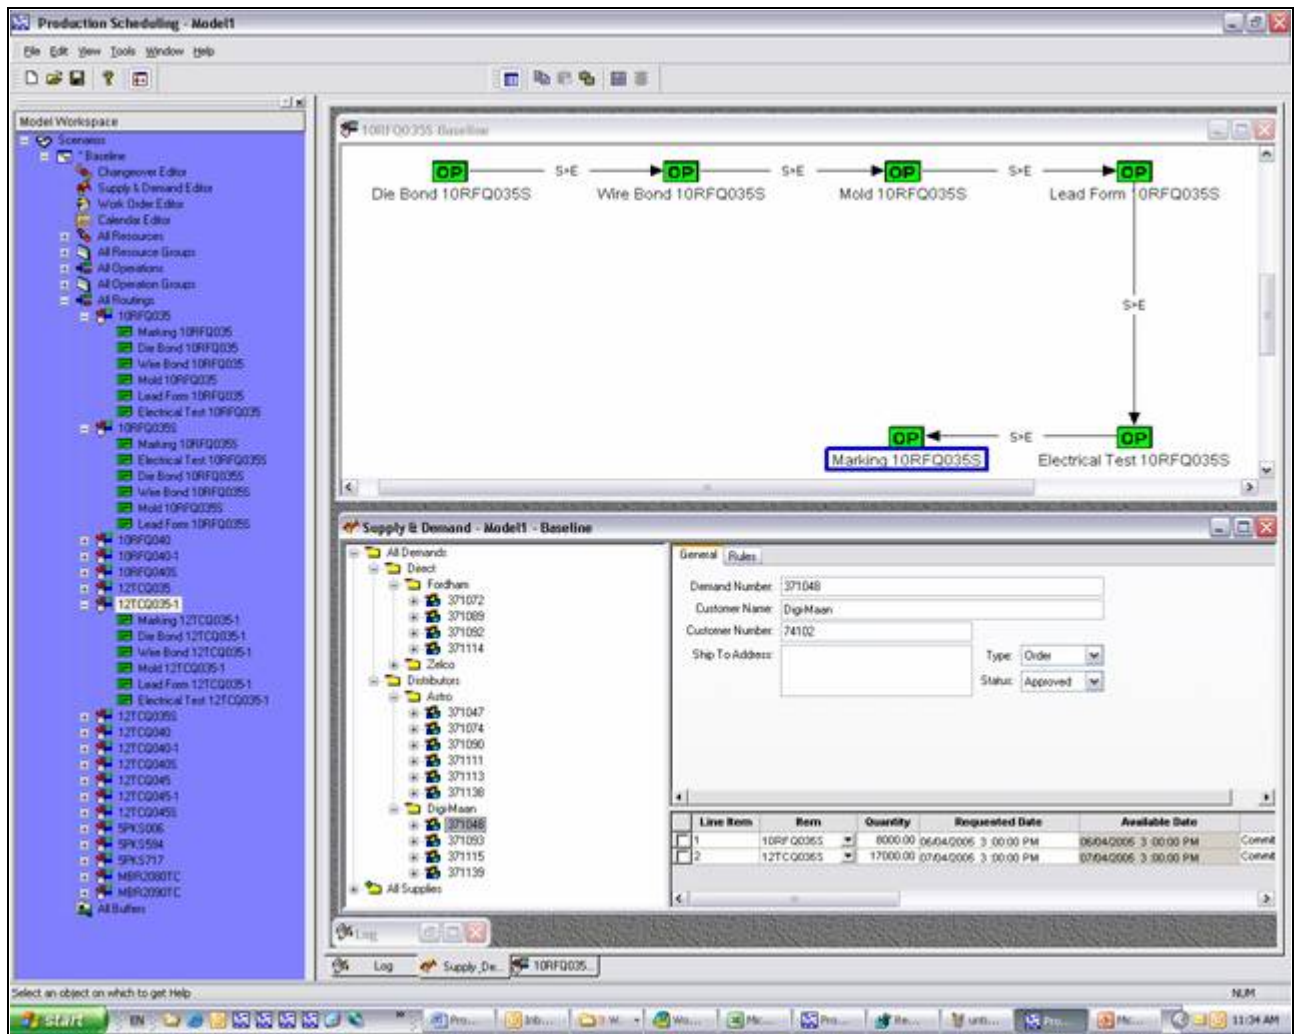
Note. The data for this example is taken from the Semiconductor model, which is shipped with the Production Scheduling application.

The flow of resources is shown below:
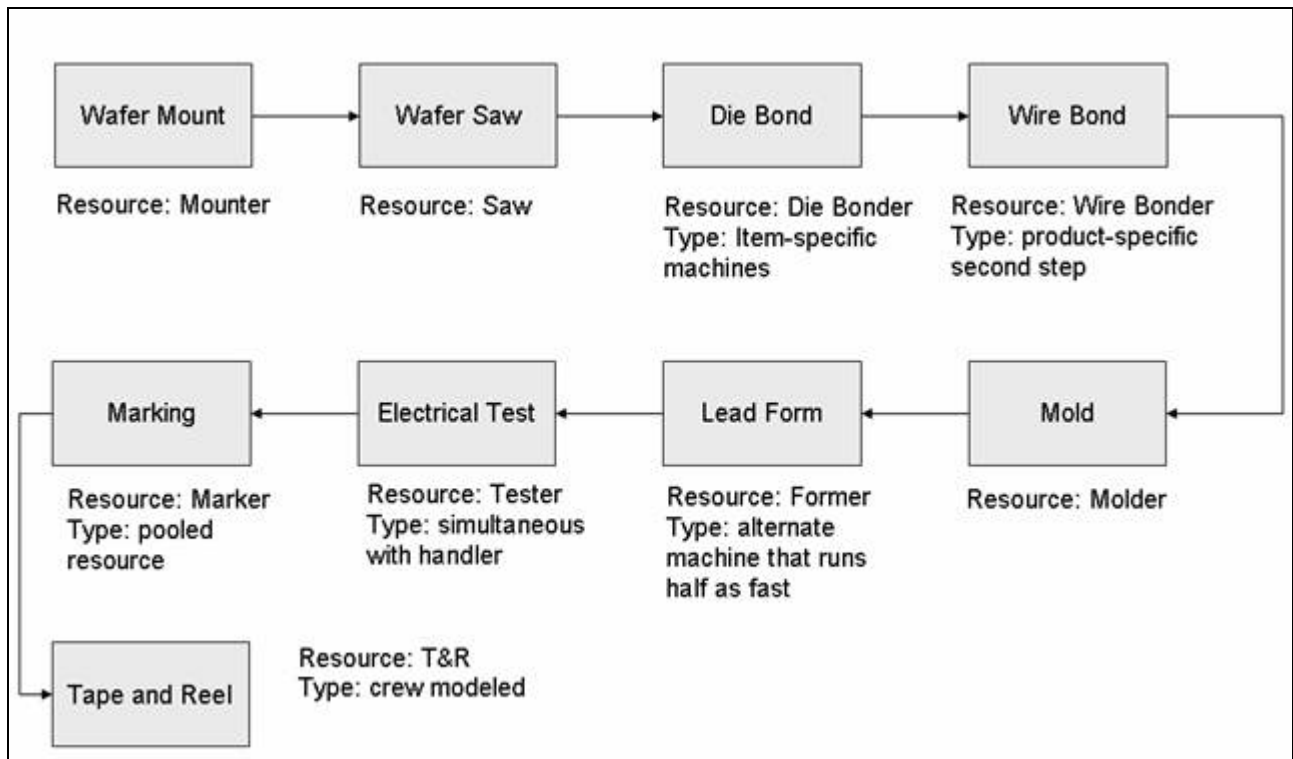
General flow of a product through a factory

The manufacturing processes are set up to resemble the routing shown in the preceding diagram. Demands have been entered for the various products.

Routing an incremental schedule

By running a solve with all constraints enabled, the demand fulfillment is really poor on a first pass schedule, as illustrated in the following screen shot.
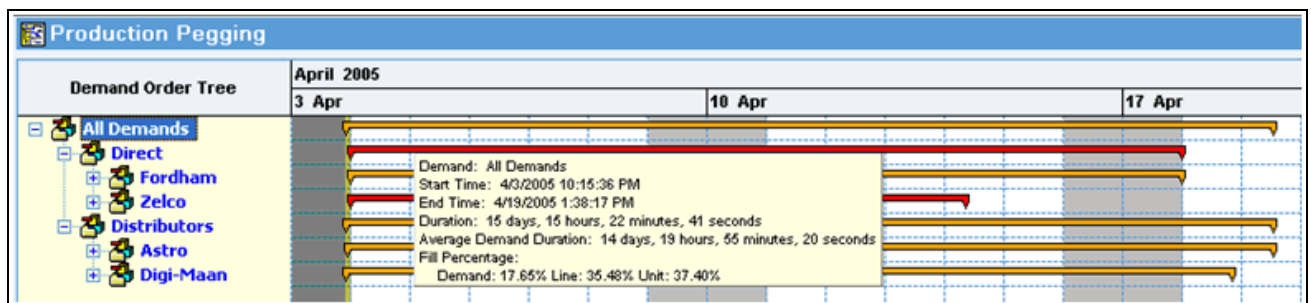
Demand fulfillment after a solve

Looking at the Resource Gantt the electrical testing machine is the dominating bottleneck resource in the factory:
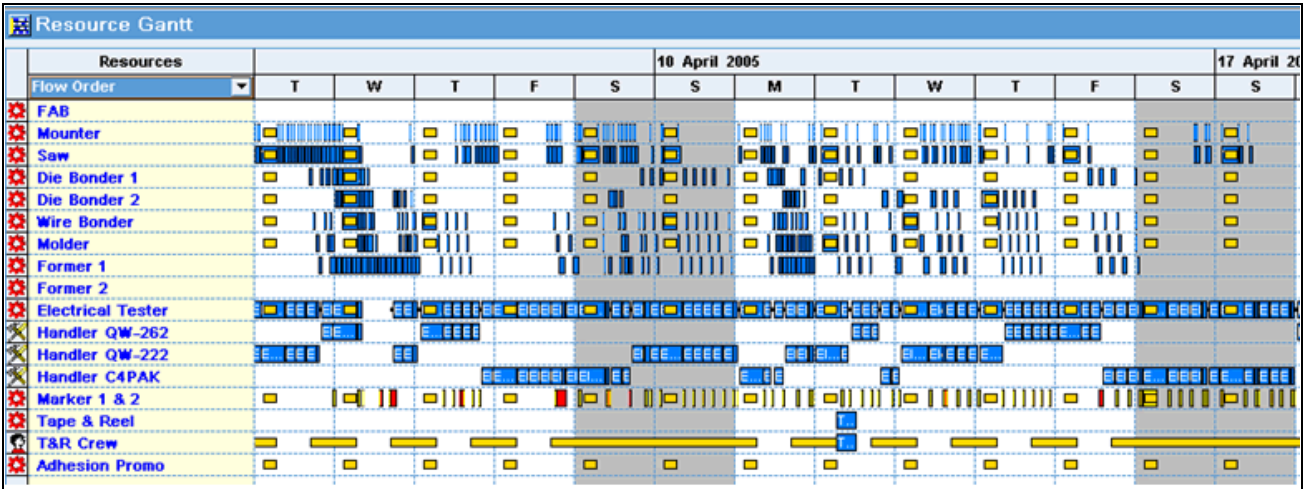
Resource Gantt

By using this feature, you can create a schedule by first solving the bottleneck resource, and then incrementally enabling upstream constraints, followed by downstream constraints.

To set up your schedule to incrementally solve for resources, you need to define the solver sequence and assign the appropriate resources to the respective stages. Create the remaining solve stages and assign the resources appropriately. In this case, the solve stages downstream resources incrementally in order of precedence, followed by working backwards, upstream, from the Electrical Test machine as illustrated.
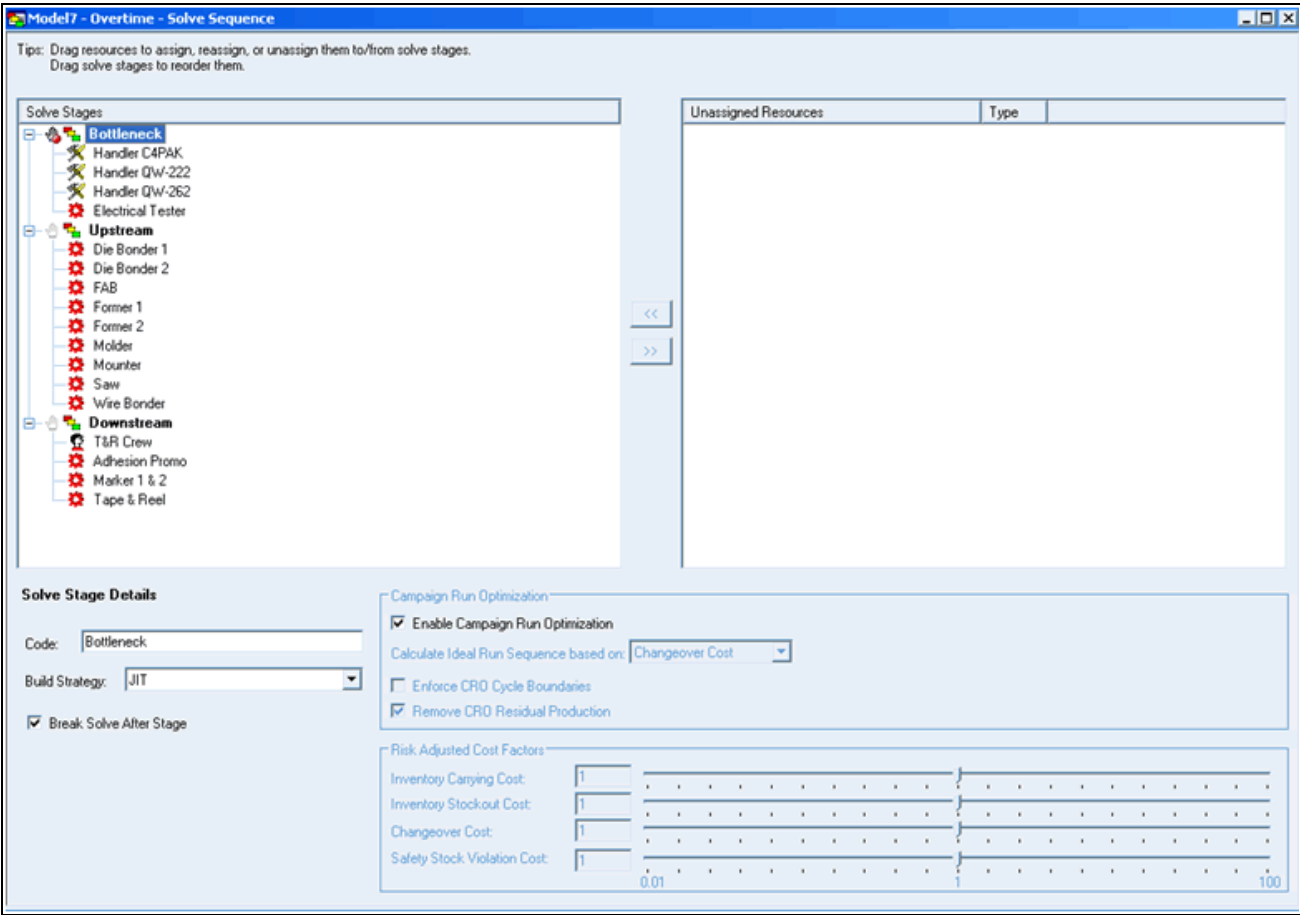
Creating the remaining solve stages

A break is inserted after the first solve stage to allow for schedule review and possible manual scheduling intervention.

Note. All resources that are not assigned to a solve stage are solved last with constraints enabled.

After re-solving with solve stages, the solve will stop at each break. From here, the schedule can be reviewed and changed by manually scheduling.
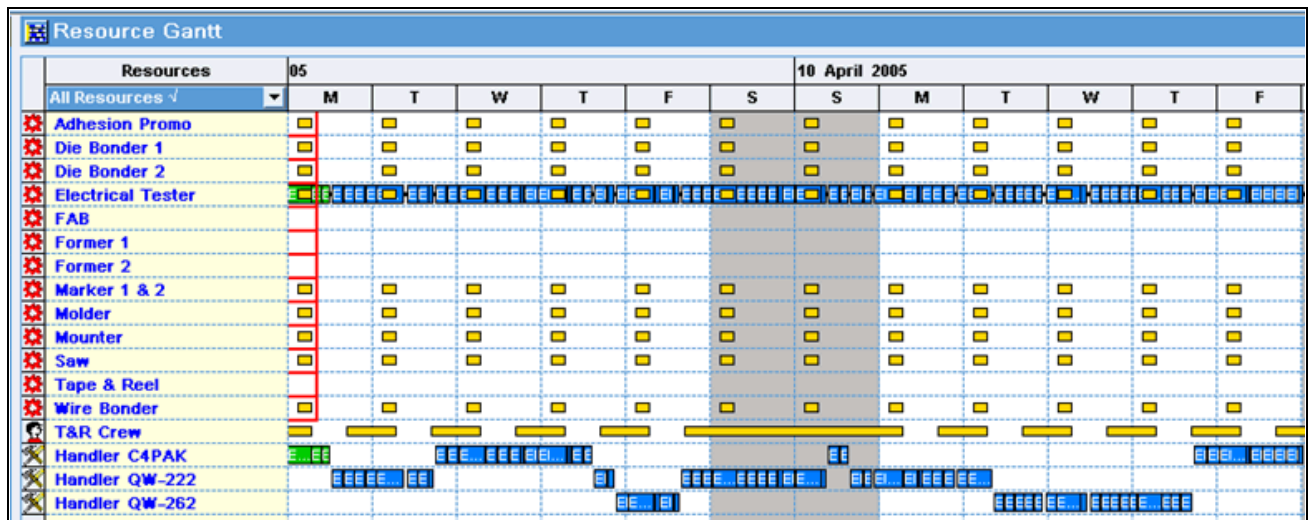
Resource Gantt review after solve break

After making a manual scheduling change, a repair solve is required to reflect the manual change after which the solve can be resumed by clicking the Resume Solve button on the toolbar where it will stop at the next solve break if defined.

After resuming the solve, the demand fulfillment has increased significantly on the first pass schedule as illustrated in the following screen shot:
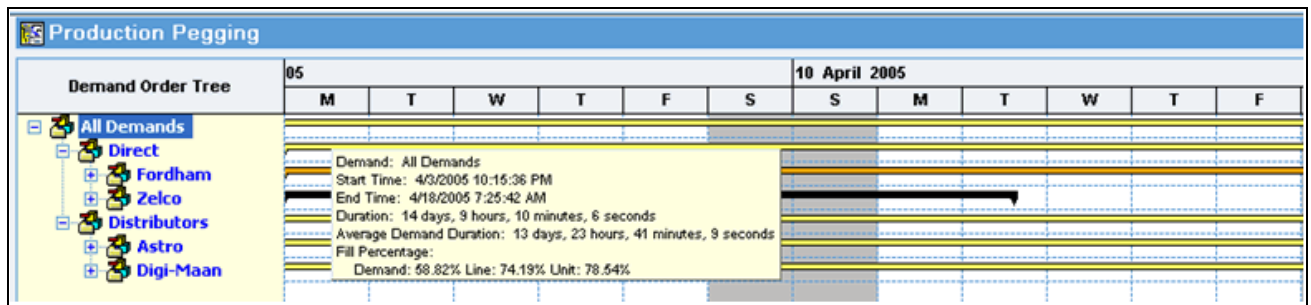
Demand fulfillment increase

However, the Electrical Test machine is still the bottleneck as shown in the Resource Gantt:
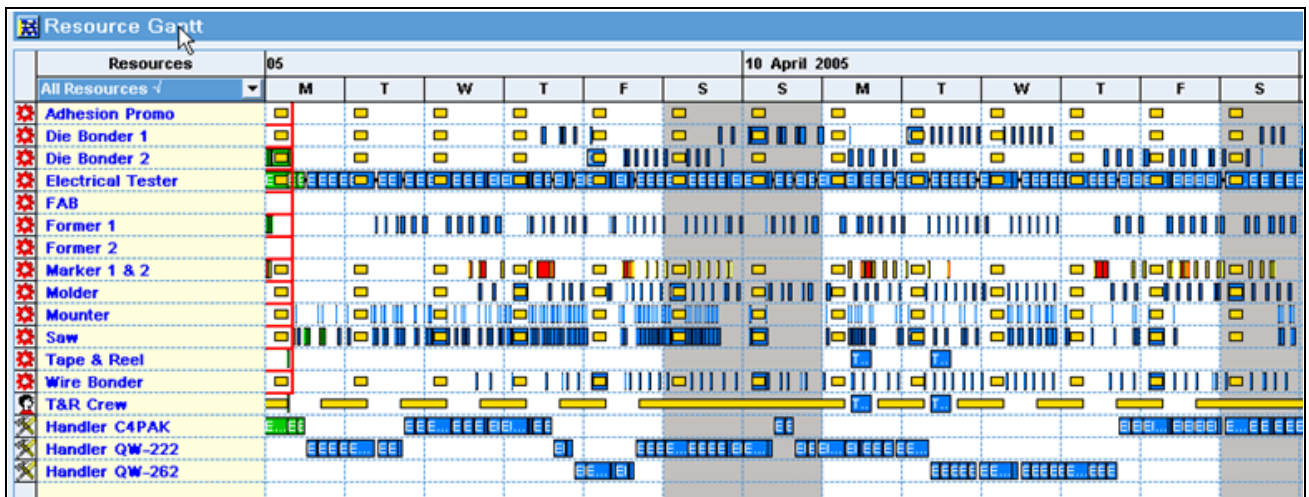
Resource bottleneck

By looking at the key performance indicators (KPIs) view, you can see the improvements from both a customer service perspective and changeover minimization perspective. In all cases, fill rates are higher using the incremental approach, total stockouts and stockout time is lower, and time lost due to changeovers is lower.

		Customer Service					Materials	Costs ($)		Manufacturing		
	Schedule	Order Fill %	Line Fill %	Unit Fill %	No. of Stock Outs	Stock Out Time	WorkOrder Fill%	Inv. Turns	Change Over Cost	Inv. Cost	Change Over Time	Machine Util. %	Throughput shifts
1	Baseline	17.65 %	35.48 %	37.40 %	15	42.72	89.50 %	431.559	1050.00	9746607.4	0.66	20.819	5381.241
2	Solve Stage	58.82 %	74.19 %	78.54 %	8	24.13	95.58 %	525.026	950.00	10574417.	0.59	20.819	5852.217

Key performance indicators

There is no black and white answer as to when the best time to use this functionality may be. Generally, if you are not satisfied with the first schedule using the generic solve logic, it may be a good idea to try incremental schedules for various solver trials.

The following are some general guidelines:

• Solve the dominating bottlenecks in the first solve stage.

• Work your way downstream by resource, keeping alternate resources in the same solve stages.

• Once you reach the end, work your way back upstream from the bottleneck resource.

The optimal sequence will result from testing various solve stage permutations and the order in which resources are solved.

Accessing the Solve Sequence Page

To access to Solve Sequence page, right-click on the Scenario from the Model Workspace and select Solve Sequence. The Solve Sequence window appears.

Creating Solve Sequence Stages

The Solve Sequence window is used to create, modify and delete Solve Stages. It allows you to assign a resource to the appropriate solve stage. This feature allows you to determine the order in which a solve is generated to resolve bottlenecks and other issues to meet your needs. The left side of the Solve Sequence page is used to define solve stages for your scenario. The right side of the page displays all resource that can be assigned to a solve stage.

To create a new solver sequence stage:

1. Access the Solve Sequence page.

2. To add a solve stage, click the Add Solve Stage from the toolbar, or right-click in the Solve Stage region and select Add Solve Stage. A new solve stage appears in the Solve Stage region.

3. To rename the stage, select a solve stage and enter the solve stage name in the Code field in the Solve Stage Details area.

4. Use the Build Strategy option to assign the solve stage as Pre-built or JIT.

5. To have the solve stopped upon completing the solve stage, select the Break Solve After Stage option. The solve is stopped at the solve stage, allowing you to review the schedule and make manual scheduling changes, if required.

6. If necessary, define the Campaign Run Optimization options for the solve stage. The Campaign Run Optimization options will be enabled if there is an applicable resource in the respective solve stage, otherwise the options are not applicable. This features allows you to define specific campaign run options to the solve stage, and these can differ from the Solve Option settings defined for the scenario.

   See "Setting Up Campaign Run Options," Configuring Solver Options, Setting Up Campaign Run Optimization, *Production Scheduling*

7. If necessary, define the Risk Adjusted Cost Factors for the solve stage. This features allows you to define specific risk factors to the solve stage, and these can differ from the Solve Option settings defined for the scenario.

   See "Risk Adjusted Costs," Configuring Solver Options, Setting Up Campaign Run Optimization, Understanding CRO Options, *Production Scheduling*

8. Create additional solve stages as needed and define their solve stage details.

9. Drag and drop an unassigned resource to a Solve Stage or perform the following:

   a. Select a solve stage.

   b. Select an unassigned resource and click <<. The resource appears under the solve stage on the Solve Stages regions.

   c. Repeat the process until you have assigned the resources to the appropriate solve stage. Resources are arranged alphabetically when they are assigned to a solve stage. The order is not indicative of a work order process.

10. If you want to move an assigned resource to a different Solve Stage, drag and drop the resource to the appropriate solve stage.

11. To remove a resource from a solve stage, select the resource and click >>. The resource is removed from the Solve Stage area and appears in the Unassigned Resource area.

Changing Solve Sequence Stages

To change solve sequence stages:

1. Access the Solve Sequence page.

2. In the left pane of the Solve Sequence page, select the solve stage that you want to move.

3. To move the solve up ahead of the previous stage or after to the next stage, click the up or down arrows at the top of the pane, or right-click on the solve stage and select the up or down arrow.

4. To move resource to a different solve stage, drag and drop the resource to the appropriate stage.

5. To remove a resource from a solve stage, select the resource and click >>, or right-click on the resource and select the Unassign option. The resource appears in the Unassigned Resources region.

6. To rename the stage, select a solve stage and enter the solve stage name in the Code field in the Solve Stage Details area.

7. Use the Build Strategy option to assign the solve stage as Pre-built or JIT.

8. To have the solve stopped upon completing the solve stage, select the Break Solve After Stage option. The solve is stopped at the solve stage, allowing you to review the schedule and make manual scheduling changes, if required.

9. If necessary, define the Campaign Run Optimization options for the solve stage. The Campaign Run Optimization options will be enabled if there is an applicable resource in the respective solve stage, otherwise the options are not applicable. This features allows you to define specific campaign run options to the solve stage, and these can differ from the Solve Option settings defined for the scenario.

    See "Setting Up Campaign Run Options," Configuring Solver Options, Setting Up Campaign Run Optimization, *Production Scheduling*

10. If necessary, define the Risk Adjusted Cost Factors for the solve stage. This features allows you to define specific risk factors to the solve stage, and these can differ from the Solve Option settings defined for the scenario.

    See "Risk Adjusted Costs," Configuring Solver Options, Setting Up Campaign Run Optimization, Understanding CRO Options, *Production Scheduling*

Removing Solve Sequence Stages

To remove a solve sequence stage:

1. Access the Solve Sequence page.

2. Select the stage that you want to delete and click the X (Delete) button at the top of the pane.

    When you delete a solve stage, all of the resources that were assigned to the stage appear in the Unassigned Resources resource.

Renaming Solve Sequence Stages

To rename a solve sequence stage:

1. Access the Solve Sequence page.

2. Select a solve stage.

3. Enter the new name in the Code field.

Setting Up Campaign Run Optimization (CRO)

This section provides an overview of campaign run optimization (CRO) and model configuration options, and discusses how to set up CRO.

Understanding Campaign Run Optimization

This section provides an overview of campaign run optimization (CRO) and provides:

* An example of CRO.

* Mid-period due dates.

CRO is an optional step to the solve that is run if you have licensed CRO in the application and have configured the data model appropriately. The primary purpose of the CRO solver algorithm in Production Scheduling is to create a production schedule that optimizes campaign run length while minimizing sequence dependent setups on a given machine. Simultaneously, the CRO solver balances the trade-off between the cost of manufacturing and carrying inventory, stocking out or violating safety stock during or across a user defined minimum manufacturing cycle time.

CRO requires separate licensing. This feature is available only if it has been enabled in the License Manager.

To use CRO, you must:

1. Configure your model to use CRO by enabling this feature in solver options.

2. Indicate which resource is campaign run optimized.

    If a resource is a member of a set within an operation, ensure that all resources are marked as campaign run optimized and that they share the same minimum cycle size.

3. Indicate the resource and operation minimum run length by time or quantity.

4. Ensure that you define item carrying costs, item stockout costs, and item safety stock violation costs.

    As a general rule, item carrying costs should always be less than item stockout costs.

5. Define machine operating costs.

6. If defining changeovers, define changeover time and cost.

An Example of CRO

This example assumes that we have three weekly manufacturing cycles. Within each cycle, there are demands for Items A, B, and C, in various quantities, which are displayed in the following diagram:


Manufacturing cycles and their demands

Each item also has a manufacturing process associated with it, which is set up as follows:

* The grinder, which is used in the production of each item, is marked as a campaign run optimized resource.

* The minimum run length is 10 units of any item when run on the grinder

* Operation lot multiple is 5 units.

* The rate is 2 hours to make 5 units of any product on the grinder

This diagram illustrates the processes for each respective item:
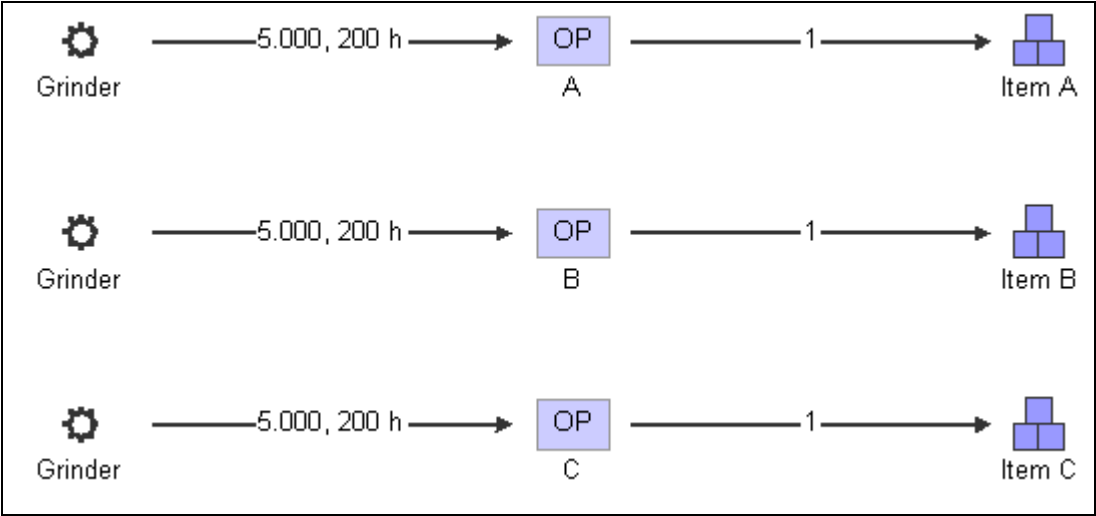
Manufacturing process for items A, B, and C

This table provides the changeover rules for this example:

Machine	From Operation	To Operation	Duration (hours)	Calendar	Cost (US$)
Grinder	A	B	1.00	Default	75.00
Grinder	B	A	1.50	Default	50.00
Grinder	B	C	4.00	Default	250.00
Grinder	A	C	2.00	Default	80.00
Grinder	C	A	4.50	Default	250.00
Grinder	C	B	3.00	Default	150.00
Grinder	*	*	1.00	Default	50.00

The CRO first determines how much of each item is needed in a given bucket, for example supply and demand netting. This diagram illustrates the ideal run length, which is calculated by aggregating demands in the user specified cycle:
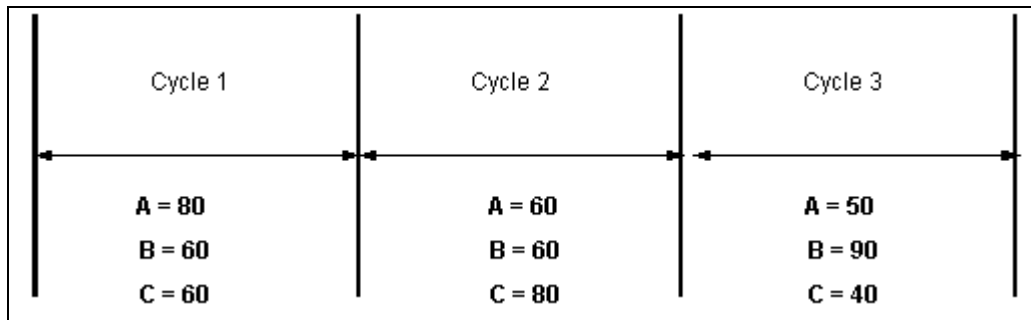
CRO run length

For example, if there are demands for 125 units for an item in cycle 1, and there is a minimum run length of 100 items, this step creates a series of operations for the same item to be 125 or more, depending on the operation lot multiple. If there is demand for only 25 units in cycle 1, the run length is calculated as 100 units.

Next, the ideal sequence within each cycle is calculated. The purpose of this step is to minimize changeovers within the user-defined manufacturing cycle times. This is based on changeover cost or time, which is user-specified.

Because the CRO algorithm works in *buckets,* the solver buckets demand for the CRO, based on the bucket boundaries. In this example the bucket boundaries are weeks. CRO then calculates a plan, based on the weekly bucketed demand, and Production Scheduling then sequences the runs according to the recommendations of the CRO.

Scenario solve results using CRO

Understanding CRO Options

For each schedule, you must first set up your Model Configuration options before you run the CRO. The options that need to be set up are:

• Minimum cycle time.

• Minimization of changeovers between cycles.

• Ideal run sequence.

• Risk adjusted costs.

• User defined ideal run sequence.

Minimum Cycle Time

Before you can run the CRO algorithm, you need to specify the manufacturing cycle time that pertains to the run length calculation. Setting a manufacturing cycle time gives the solver insight into your product manufacturing cycle and enables it to determine the appropriate run size that is required to optimize product resources. Generally, the run pattern repeats across cycles.

Minimum Run Cycle is user defined, for example one day, two weeks, and so on. This can be a global option or it can be resource specific.

Minimization of Changeovers Between Cycles

Minimizing changeovers between cycles avoids changeovers between the last operations of the previous manufacturing cycle and the first operation of the of the following cycle. The solver always calculates the optimal sequence based on changeover durations or cost. The ideal sequence is consistent for each cycle, although Production Scheduling won't always run items in the sequence due to supply and demand variability or pull-forward production.

Minimizing changeovers between cycles enables the solver to weigh the changeover cost between cycles against the carrying cost of an item in the next cycle. Production of the item may be pulled forward into an earlier cycle to eliminate a subsequent changeover.

Ideal Run Sequence

Ideal run sequence is the order in which events happen within a cycle.

If changeovers are specified by cost, it may result in changeovers which are slightly longer, but cost optimal. If changeovers are specified by time, the time spent on changeovers is minimized but the costs may be higher.

Note. The costs you specify in your scenario greatly impact the quality of your schedule. Make sure that you have accurately specified the costs of inventory carrying, changeovers, safety stock violation, and inventory stockout. Costs are normally calculated during model prototyping.


Risk Adjusted Costs

When the CRO algorithm is run, it looks at various costs in the model, including inventory carrying, inventory stockout, changeover, and safety stock violation costs when it attempts to create an optimal production schedule. These costs are typically imported from your ERP system.

The Risk Adjusted Costs functionality in Production Scheduling enables you to quickly scale the relative costs of inventory carrying, inventory stockout, changeover, and safety stock violation costs. The actual costs in the model do not change, but they are "scaled" in the Production Scheduling solver when it makes trade-off decisions in creating a schedule.

The cost factors default to a value of 1. You can quickly and easily change the relative importance of the specific cost by either moving the slider bar to the left or right of 1, or by manually typing in the value in the field beside the slider bar. The acceptable values for each of the respective scale factors range from *0.01* to *100.*

User Defined Ideal Run Sequence

In many environments, you may simply wish to dictate the order of changeovers to run. You can specify the sequence operations should be scheduled in CRO according to attribute (available for CRO resources only). This is an alternate sequencing strategy compared to the current CRO sequencing strategies of Cost or Time.
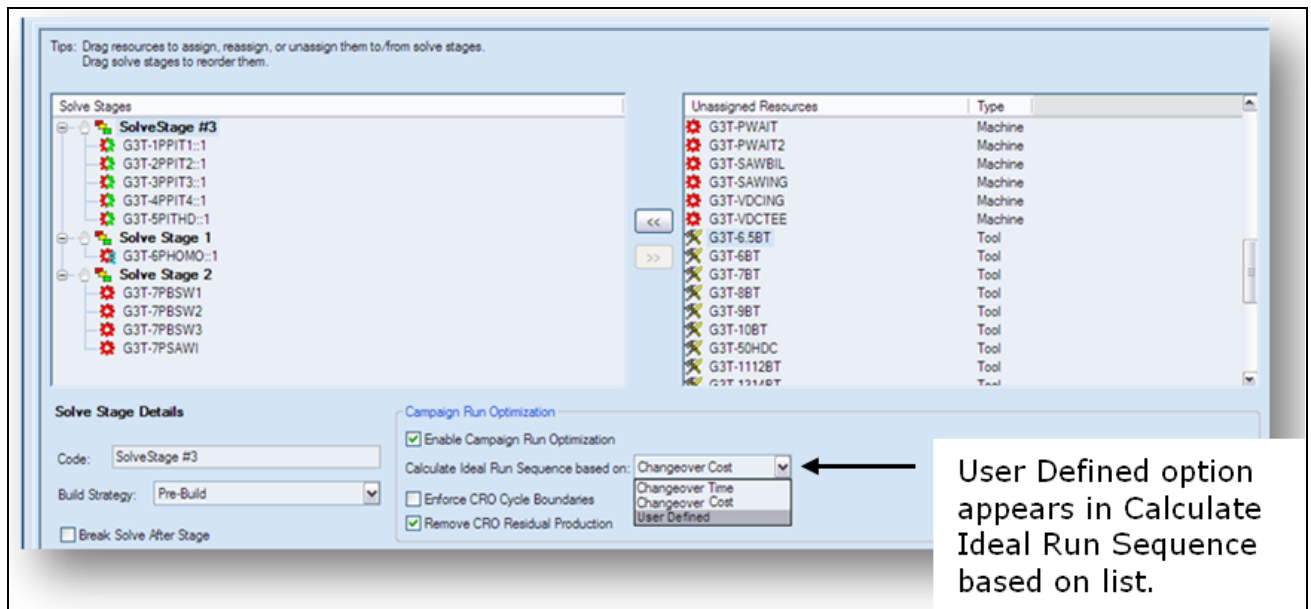
Example of User Defined option in Calculate Ideal Run Sequence list

The mechanism in which the "ideal" sequence is defined is by selecting and sorting the associated attributes and their attribute values on CRO resources in the model.

For example, take attributes "Size" and "Grade". Within "Size" there are 1/4 Inch, 1/3 Inch and 1/4 Inch as attribute values. Within the "Grade" attribute, there are Grades A and B.

Assume the user defined Ideal Sequence is to cycle by size and grade within the size, as outlined in the following example:

1/4 Inch-Grade A -> 1/4 Inch Grade B -> 1/3 Inch-Grade A -> 1/3 Inch Grade B -> 1/2 Inch-Grade A ->1/2 Inch Grade B

This would be defined in the Attribute Processing tab of Production Scheduling where attributes are associated to resources.
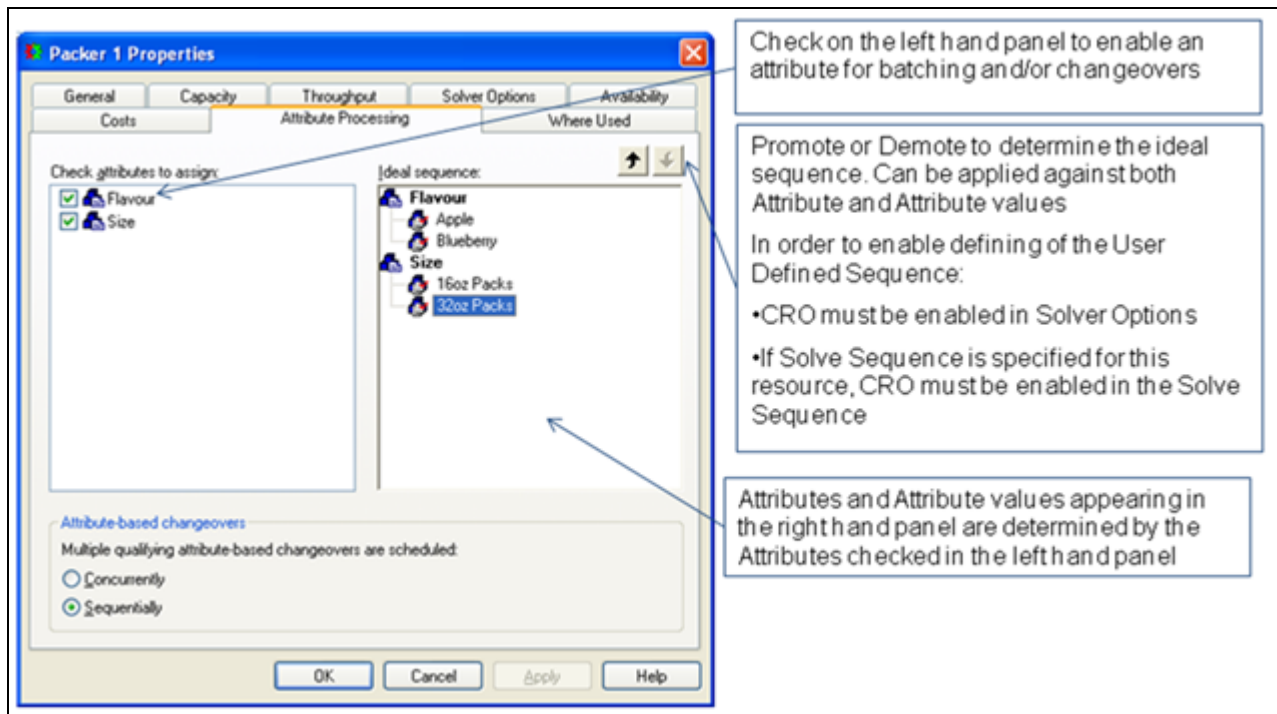
Example of Attribute Processing tab where user defined ideal sequence is defined

Essentially, the order in which the selected attributes and attribute values in the Ideal sequence column are structured dictates the ideal sequence of operations (the operations must be associated to these attributes). For example, if Grade was more important than Size, then you could move the Grade above the Size attribute using the Promote or Demote button above the Ideal sequence area. The sequence would then cycle by Grade and then Size within the Grade.

Model Building Considerations

When building models, you should consider certain topics:

*   Optimizing single-capacity machines.

*   Late demand.

*   Campaign run optimized machines.

Only single-capacity duration machines can be optimized, and no more than one machine per routing can be optimized for repetitive manufacturing. If the machine is part of a resource set, then all machines within the set must be optimized for repetitive manufacturing.

Because the CRO algorithm is based on buckets, it is possible that the algorithm may not recognize demand if a given demand is late within a given bucket. For example, if demand is due at noon within a bucket, the CRO is solving on the basis of daily buckets, and the sequence of operations determines yields on available time of 4:00 p.m., the CRO phase of the solve does not see this demand as late.

Because the sequencing phase of the solver cannot change the operation sequence determined by the CRO, the demand may appear late within a given cycle. If inventory levels are below zero on a cycle boundary, then the CRO sees an item stockout. However, there is logic in the solver that internally sets the due date of the late demand to a previous bucket if there is sufficient idle time on the campaign run optimized machine. The CRO is then internally run again.

Within the manufacturing flow of an item, from raw materials to finished product, you can specify either one resource or resources in a resource set that can be campaign run optimized. This diagram illustrates run optimization on a single resource, which is Machine 5:
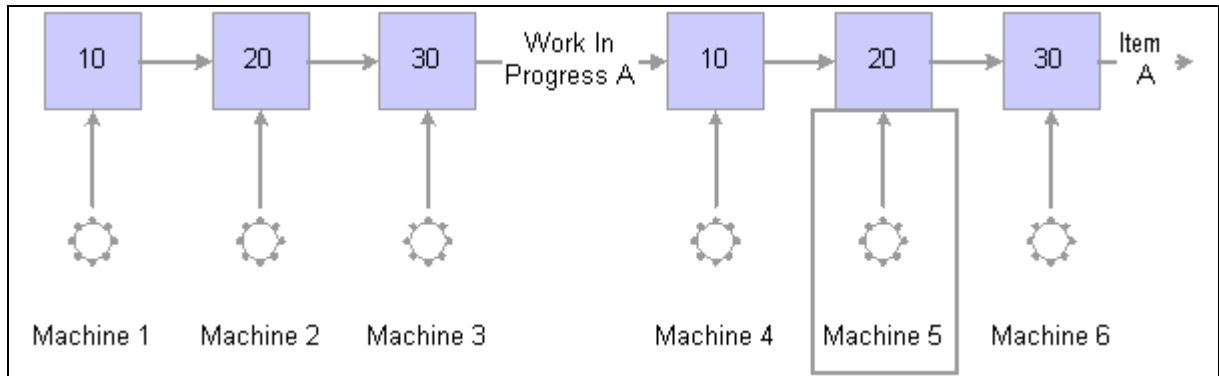
Campaign run optimization with one resource optimized

The resource or resources in a resource set that you designate as campaign run optimized may be used in several different operations if the operations are not producing the same final item. This diagram illustrates this concept for Machine 5:
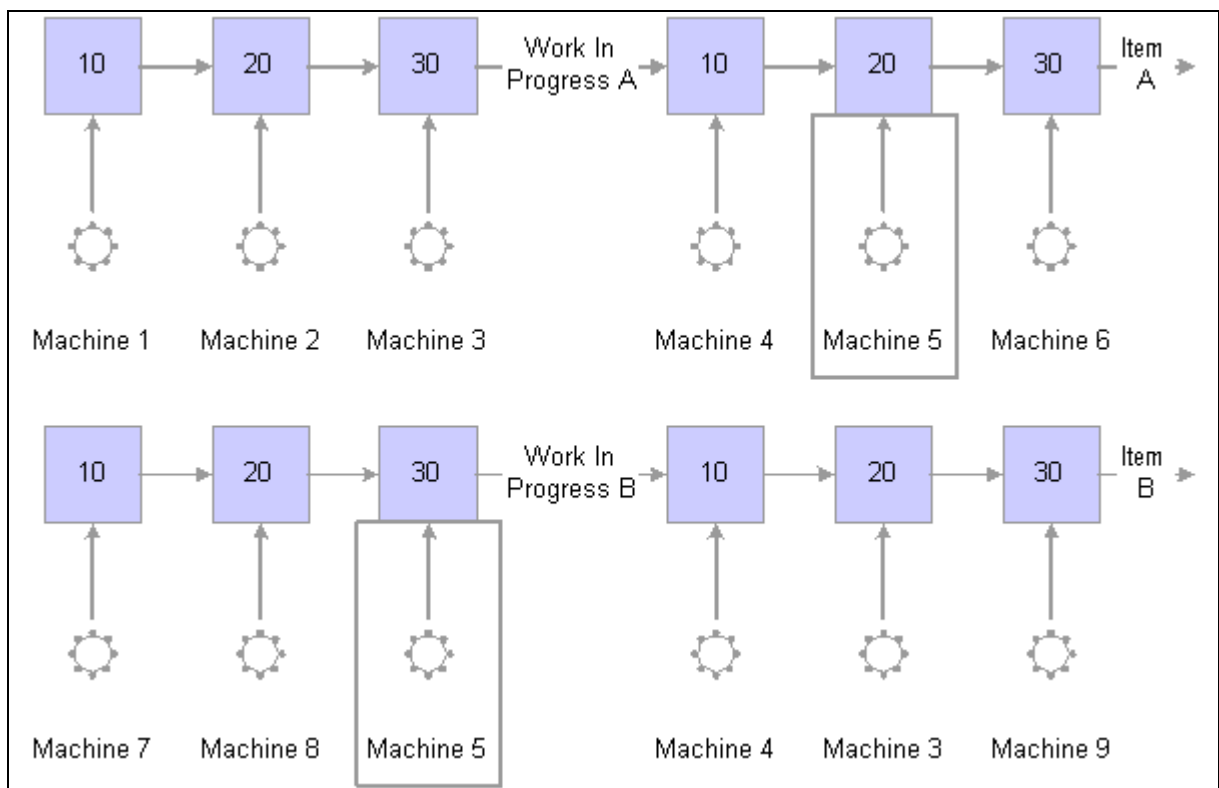
Campaign run optimization with one resource campaign run optimized for different end products

**Note.** Only single capacity and batchable machines can be designated as campaign run optimized.

A campaign run optimized resource cannot be set to relaxed.

Some scenarios, which are explained in the following paragraphs, are not conducive to run optimization:

1. You cannot specify a resource as campaign run optimized if the machine is used in a material flow that feeds into another material flow, which has a campaign run optimized machine.

2. A resource can be used only once in the entire material flow.

    Multiple runs on the same resource in the same material flow presents conflicting objectives.

The following diagrams illustrate these restrictions on run optimization. This diagram illustrates a scenario in which Machine 3 is campaign run optimized, but feeds into a material flow in which Machine 5 is also campaign run optimized:
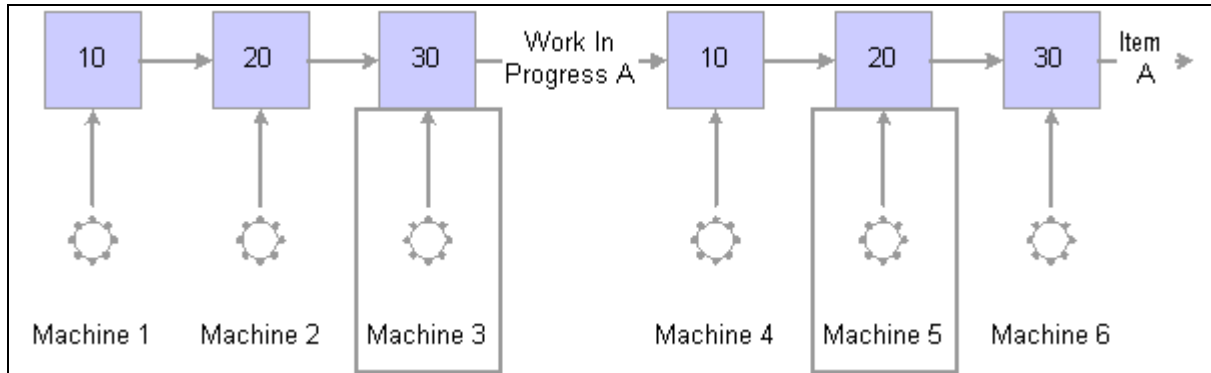
Two campaign run optimized machines with the same finished product

This diagram illustrates one resource being campaign run optimized more than once in the same material flow:
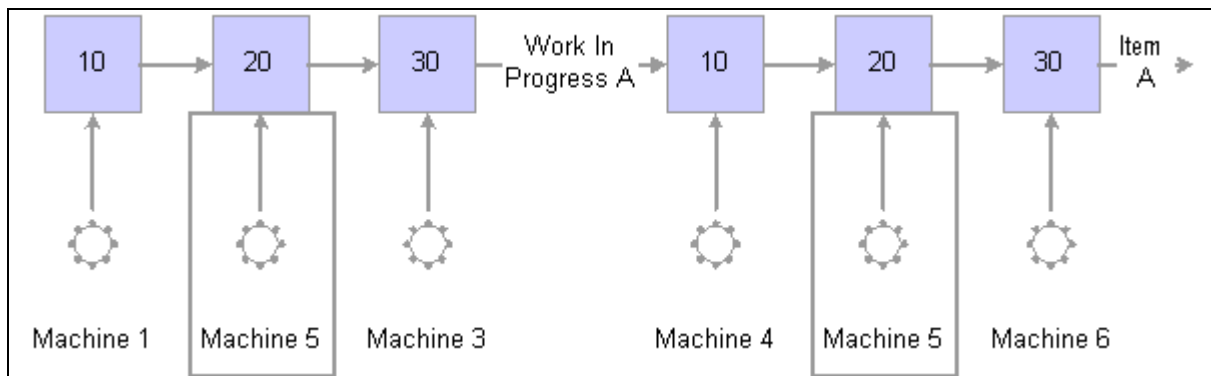
One resource set twice to run optimization

Minimum Run Length Specification

In Production Scheduling, you can specify a minimum run length. The minimum run length can be set in units or in time. If you set a minimum run quantity on the resource, the operation, which is producing the item, must run a minimum quantity of the item. Similarly, if you set a minimum run in terms of time, the operation must continue for the specified amount of time.

See "Creating an Operation Diagram," Creating Production Scheduling Models, Defining Operations, *Production Scheduling*

Multi-Stage Campaign Run Optimization Limitations

When using multi-stage campaign run optimization, it is important to note the following limitations:

- If multi-stage is setup (per item material flow – it can still be single per flow without a solve stage as per existing functionality) but solve stages do not exist, then CRO will be disabled and a warning placed in the solver log. This is a configuration error (for example, solve using the CSP). Validation should say "Multi-stage Campaign Run Optimization is configured without solve stages. CRO will not be utilized for this schedule"

- Multi-stage CRO within the SAME routing is NOT supported. CRO resources called in a multi-stage fashion must be in different stages.

- It is best to limit the number of CRO stage. Keep CRO stages to a minimum.

- If alternate resources or alternate routings for the same item are present, they must exist in the same solve stage.

- Resources can not recur across solve stages. For example, users can not solve the injection stage more than once.

Setting Up Campaign Run Options

To set up campaign run options:

1. Right-click a scenario, and select Solve Options.

2. Select the Campaign Run Options tab.

3. Complete these fields:

Enable Campaign Run Optimization	Select to indicate that you want to use this feature. When you select this option, the remaining fields on this page are available.
Minimum Cycle Time	Enter a value and select the time measurement required to cycle through demand for the items that are produced. Time measurement options are *Shift, Days,* and *Weeks.*
Resource Specific	Select to indicate that minimum cycle time is different for each resource.
	When this option is selected, set the cycle time and select the Minimize Changeovers between cycles check box on the resource.
Minimize Changeovers between cycles	Select to indicate that you want Production Scheduling to consider what operations ran in the previous cycle. If possible, the application schedules an operation that can run on this resource without a changeover.
Calculate Ideal Run Sequence based on	Select *Changeover Cost* if you want to minimize the changeover cost in each cycle.
	Select *Changeover Time* if you want to minimize the changeover time in each cycle.
	Note. To accurately minimize changeover costs, you must enter a cost value for each changeover rule.
	Select *User Defined* if you want to define the sequence order you define as the criteria used to solve the model.
Enforce CRO Cycle Boundaries	Select this option to enforce the CRO cycle boundaries.
Remove CRO Residual Production	Select this option to remove CRO residual production.
View Campaign Run Optimization solution results only	Select to view the CRO solve results only. This helps to understand which decisions have been made by the CRO algorithm prior to them being fed into the base solver. This flag is generally only used in model prototyping and not set within a production environment. This flag is not persisted in the data model and is session specific.
Risk Adjusted Costs	Use the scale to select the appropriate scale factor.
	See "Risk Adjusted Costs," Configuring Solver Options, Setting Up Campaign Run Optimization, Understanding CRO Options, *Production Scheduling*

Chapter 9

Evaluating Production Schedules

This chapter provides overviews of production schedule evaluation and schedule views and discusses how to:

- Analyze scheduled resources.

- Analyze scheduled items.

- Analyze scheduled operations.

- Use combined views.

- Work with alerts.

- Work with key performance indicators (KPIs).

Understanding Production Schedule Evaluation

This section discusses:

- Solve status.

- Solver statistics.

- Solver log data.

Solve Status

When the Production Scheduling is solving a schedule, a number of messages appear in the status bar to indicate what tasks the solver is performing at that moment. Real-time messages enable you to keep track of each phase of the solve process as it occurs. The application produces the following messages while solving a production schedule:

Solver Message	Expanded Definition
Preparing Model	The model is being prepared for the solver.
Calculating Net Demand	If the Campaign Run Optimization (CRO) algorithm is enabled, it first calculates net production requirements on the CRO machine based on the supply and demand picture in the model.

Solver Message	Expanded Definition
Optimizing Run Length	If the CRO algorithm is enabled, it optimizes sequence dependent setups, run length, and simultaneously balances the trade-off between building and carrying inventory, stocking out and, violating safety stock levels.
Solving - Assigning Operations to Resources	Based on the model information, the solver is creating all of the operations that are required to meet the model demand and assigning them to applicable resources.
Solving - Sequencing	In this phase, the solver is sequencing the operations while respecting all of the constraints and trying to meet demand due dates.
Preparing Schedule	In this phase, the solver data is converted into a schedule.
Refreshing Views	The solver is updating any open schedule views to reflect the current schedule.
Scheduled in <x seconds/minutes>	The schedule has been solved in the specified time.
Schedule Infeasible	This message appears if the schedule is infeasible. Check the log window for details.
Solver Stopped	The user has stopped the solver.

Solver Statistics

Production Scheduling creates a solver statistics file that captures important solver information after solving every schedule. If you have selected the Log Messages option in the Options dialog box, then solver statistics are written to a SolverStatistics.txt file every time that you solve a schedule. This file can be viewed directly from the Log window.

You can examine the solver statistics file to obtain more information about the following solve characteristics:

• Determine the solver performance for the solved schedule.

• View demand details for the solved schedule including demand fill percentage, demand total, and on-time and late demand statistics.

• View operation and resource totals for the solved schedule.

• View iterative and heuristic information about the solve.

Solver Log Data

The solver log provides information about the number of demands that were met on time within each solve that is performed by Production Scheduling during the current session. By viewing these solver statistics in the Log window, you can determine how successful the solve was.

When the schedule is infeasible, the system displays one of the following messages in the log:

Solver Message	*Explanation*
Insufficient Resource Capacity	Not enough crew, machine, or tool capacity is available for the specified resource.
No Producing Operation in Model	All items in the model require a source of supply whether it come from a supplier or a manufacturing operation.
Insufficient Item Maximum Storage Capacity	Too much inventory is being produced to fit in the allocated storage space.
Routing Precedence Constraint Error	The solver has encountered an issue with an operation sequence and can not resolve the conflict. For example, Operation B might be firmed to occur before Operation A but it can not feed it sufficient inventory.
Could not Compute Ideal Operation Sequence for Campaign Run Optimization	The solver was unable to generate a sequence of operations based on your choice of using changeover duration or cost.
Unable to Initialize Campaign Run Optimization	An internal or licensing error has prevented the Campaign Run Optimization from initializing.
Specified Minimum Cycle Time too Short for Campaign Run Optimization	An operation's minimum run length parameters cannot fit within one cycle. Decrease the minimum run length.

Note. If you have selected the Log Messages option in the Options dialog box, then these statistics are saved to a SolverStatistics.txt file each time that you solve a schedule. This file is written to the same location as your log file.

Understanding Schedule Views

This section discusses:

- Schedule views.

- The Production Pegging view.

- The Resource Gantt view.

- The Item Graph view.

- The Item Gantt view.

- The Operation Gantt view.

- The Resource Utilization view.

- The Resource Contention view.

- The Mulit-Capacity Resource Graph view.

- Combined views.

Schedule Views

From a number of schedule views, schedules can be improved using manual scheduling and repairing techniques. Schedules can be compared in the Key Performance Indicators window against a number of objective measures to determine the best schedule. Schedule Views is the collective term for Gantt diagrams and graphs in Production Scheduling. The workspace in which these Gantt diagrams and graphs appear is known as Schedule view. You can view problems with the schedule by browsing the alerts in the Alerts grid.

This section discusses:

- Schedule view groups

- Operations

- Changeovers

- Merged operations

Schedule View Groups

Groups enable multiple operations and resources to be combined as a group when viewed in the model workspace and schedule views. Organizing resources and operations into logical groupings can help you organize your model data. You can create a default resource group for each schedule that appears when views are first opened.

All schedule views, except the Production Pegging view, have a drop-down menu that lists the defined groups. When you select a group from the menu, the schedule view changes to display only those resources or items belonging to the group. Each schedule view can have its own default group view. The current default group in each view appears with a check mark beside its name.

Operations

When using some of the schedule views, information is available about the scheduled operations to assist you when evaluating your schedule. In addition, the run times, duration, and assigned resources (if a resource set is used) can be changed directly from schedule views.

You can view information about all operations from the Demand, Operation and Resource Gantt views. Only information about single-capacity operations is available from the Resource Gantt view.

The Operation Properties window contains the following tabs:

Tab Name	Description
General	This tab contains the operation name, code, and the demand and work order that it is fulfilling. From this tab you can launch the Work Order and Supply and Demand editors to view and edit the associated demand and work orders. You can also access the operation and routing using the icons at the top.

Tab Name	Description
Times	This tab contains the selected operation's start and end times, earliest start and latest end times, duration, run, and delay. This tab also indicates whether the operation has been set to run at a specific time.
Resources	This tab contains all of the resources that are associated with this operation, including their quantities. Item resources are displayed as negative when they are consumed, because they are not renewable.
Document	This tab lists associated files and the location of the files.

A quick summary of an operation can be viewed for single operations. Details including the name, start time, and duration of the operation are included in this quick summary. A quick summary for an operation is available in the Demand, Resource, and Operation Gantt views.

Changeovers

The transitional time between one operation running on a specific machine and another operation is called a changeover. Changeovers appear in schedule views as solid black lines between operations.

Changeover events that affect resources appear in the Production Pegging, Resource Gantt, and Resource Utilization views, according to the rules set in the Changeover editor.

Merged Operations

To help simplify your schedule views, adjacent operation instances of identical types can be represented as one production run. Within the Resource and Operation Gantt views, you can view the schedule with each operation instance displayed separately or as a single operation. This option is toggled with the Merge Operation button. Click the Merge Operation button again to display the operations separately.

As a result, manual scheduling is simplified. Instead of reassigning each batch in a run, you only need to reassign the total production run by using the Cut and Paste buttons. Properties for merged operations can be viewed, but they cannot be modified. The properties for merged items display the aggregated properties for each discrete operation in the merged operation. Pegging can also be performed on merged operations.

Once operations have been merged, you can view information about these consolidated operations from the Operation and Resource Gantt views. The Resource Gantt views only provide merged information about single operations, whereas the Operation Gantt displays information about all the operations.

The Production Pegging View

This section discusses:

• The Production Pegging view.

• The order tree.

• Makespan bars.

• The Production Pegging - Demand Centric view.

- The Production Pegging - Work Order Centric view.

- Operations.

- Resources.

- Supply events.

- Partially filled demand.

- Purchase recommendations.

The Production Pegging View

In the Production Pegging view, you can display several important elements of your schedule in a simple, easy to use format. Among the many aspects that you can access are:

- Multistage pegging of work orders

- Resource requirements

- Consumed items

- Net new production recommendations

- Purchase order recommendations

Production Pegging lets you view production from a demand point of view (demand centric), or a work order point of view (work order centric), which brings up a Work Order Gantt chart. To alternate between the demand centric and work order centric views, click the View Work Order Pegging or the View Demand Pegging toggle button on the menu bar. This toggle button has descriptive text that changes depending on the view that is displayed.

The Order Tree

When open to the demand centric view, this order tree displays all of the scheduled demands as listed in the Supply and Demand editor. You can expand the tree to display the line items and operations that are associated with each work demand order.

When open to the work order centric view, the tree displays all the scheduled work orders as listed in the Work Order editor. You can expand the tree to display the details of operations and routings that are used by each work order, and also see pegged upstream work orders. In addition, you can look at the production recommendations and purchase recommendations.

The makespan for each folder and order appears in the time grid, similar to other schedule views. The makespan is represented by a bar with a diamond at each end to indicate the total amount of time that is required to fill an order. This bar appears in a contrasting color to provide information about the percentage of the orders filled either in demand orders or folders.

Makespan Bars

In the Production Pegging view the makespan bars are:

- Folder based.

- Demand based.

- Line item based.

- Routings based.

The makespan bar has a different meaning based on its color.

The makespan and the order of each Work Order Centric or Demand Centric folder in Production Pegging are displayed in the time grid, similar to other schedule views. The makespan is represented by a bar with a diamond at each end to indicate the scheduled beginning and end of manufacturing for individual demand orders and their associated demand folders. The makespan bar is colored to provide information about the percentage of the orders filled in either orders or folders.

Type	Significance
Solid black bar	100 percent of orders filled, based on unit fill.
Solid yellow bar	Between 66 and 100 percent of orders filled, based on unit fill.
Solid orange bar	Between 33 and 66 percent of orders filled, based on unit fill.
Solid red bar	Less than 33 percent of orders filled, based on unit fill.

Pop-up windows provide information about each work order and folder when the cursor pauses on its makespan bar. Folder makespan bars represent aggregate information about contained work orders, and work orders summarize line item information. The following data is displayed in pop-up windows for demand and work orders in Production Pegging:

Field	Description
Demand (demand folders only)	Name of the demand folder.
Demand Number (demand orders only)	Name of the demand order.
Start time	Start time for the first demand order in the demand order or folder.
End time	End time for the latest ending demand order in the demand order or folder.
Duration	Length of time between the start time for the earliest demand order or line item and the latest ending demand order or line item.
Average Demand Duration (demand folders only)	Average makespan for demands within the folder.
Average Line Duration (demand orders only)	Average makespan for line items within the demand order.

Field	Description
Fill Percentage	Demand (only available for folders). Percentage of the orders that are filled in the selected folder.
	Line. Percentage of the item lines that are filled within the order or in the folder.
	Unit. Percentage of the units that are filled within the order or in the folder.
Instance Count (routings only)	The number of instances of this routing in this work order.

The following data is displayed in pop-up windows for work orders in Production Pegging:

Work Order Centric View	
Field	**Description**
Work Order	Type of work order. For example, Production , Maintenance, Finished Goods, or Distributors.
Start Time	Start date and time of the first operation in the routing or folder.
End Time	The end date for the operation.
Average Duration	Average makespan for orders within the folder.
Unit Fill Percentage	Percentage of the units that are filled within the order or in the folder.
Requested Date	The date on which the order is requested.
Available date	The date on which the order is available.
Duration	The duration of the work order.
Pegged Demands (for work order only)	The demands that are pegged to this work order.

The makespan of each line item and routing is displayed on a time grid. Similar to the makespan for demand orders and folders, the makespan is represented by a bar with diamonds at each end, symbolizing the start and finish of the line item or routing. If the demand line item is running late or is available early, the bar is partially displayed in color. Each line item can be expanded to display the routings and operations that are involved, as well as the resource details and supply events.

Purchase order recommendations are represented by solid green bars with a trailing diamond. The bars for purchase order recommendations represent the lead time required to procure the purchased item. The diamond represents the date when the items in the purchase order become available to use in a routing.

An additional diamond is included on the bar to indicate either the requested date for line orders scheduled to be available late, or the available date for line items scheduled to be available earlier than requested.

In addition to the third diamond, the makespan bar is displayed in color between the second and third diamonds, indicating whether the line item is available earlier or later than its requested date, and by how much.

The following table summarizes the different types of line item makespan bars that are displayed in the Production Pegging:

Type	Significance
Solid black bar	Line item scheduled to be manufactured on time.
Black and red bar	Line item scheduled to be available after the requested date.
Black and green bar	Line item scheduled to be available before the requested date.
Solid red bar	Line item to begin manufacturing after the requested date.
Solid green bar	Line item scheduled to be early.
Green diamond	Line item is procured from a supplier and ready to be used in a routing.

Pop-up windows display the following information about each demand order when you place your cursor over a specific demand line item makespan bar:

Demand Number	Name of the demand order.
Latest Requested Date	Date when the order has been requested.
Latest Available Date	Date when the order will be available.
Start Time	Start time for the line item.
Duration	Length of time between the start time and the end time.
Average Line Duration	The average duration of one line item.

Routings are displayed on the time grid by makespan bars. The color of these makespan bars is based on whether the routing is on time or late.

Type	Significance
Solid black bar	Routing is on time.
Solid red bar	Routing contains critical bottleneck operations.

Although subroutings are not displayed with a makespan bar, their operations are displayed on the Production Pegging view.

A pop-up box displays the following information:

Type	Significance
Routing	Routing name.
Start time	Start time for the routing.
End time	End time for the routing.
Duration	Length of time between the start time and the end time.

Production Pegging - Demand Centric View

The Production Pegging - Demand Centric view enables you to visually display the start times, end times, and duration for each demand, along with its associated line items and routings. From this view, you can see how complete the orders in a folder are, how complete a specific order is, and whether scheduled line items are early, on time, or late. The system distinguishes critical routings and operations. The system does not display demand orders on hold.

Note. For the Production Pegging demand centric view to properly function, you must specify the primary operation in routings and primary output items in operations. This pegging information is essential to display the sequence of operations and routings that are involved in the production of any order line item.

The Production Pegging - Demand Centric toolbar provides a number of buttons that are available only in Production Pegging. From this specialized toolbar, you can filter demand orders according to their criticality to assist you in analyzing the schedule's handling of demands. You can filter the demand orders to list only orders of a specific type, punctuality, and magnitude of punctuality. The Production Pegging retains the last filtering settings used for future analysis.

The following criteria are used to filter demand orders:

Filter	Name	Description
Primary	Order Type filter	Enables the demand orders to be filtered by the type of demand. Options include all, orders, forecasts, and safety.
Secondary	Order Punctuality filter	Enables the demand orders to be filtered by punctuality. Options include all, early, on time, and late.
Tertiary	Order Magnitude Filter	Enables the user to specify the magnitude of punctuality when filtering. By default, these options include all, >12 hours, >1 day, >2 days, and >3 days.

Note. The number of late orders is shown in the status bar at the bottom of the screen when you use the filtering in the view.

The Production Pegging - Work Order Centric View

The Production Pegging - Work Order centric view enables you to visually display the start times, end times and duration for each work order , along with its associated operations and any pegged work orders. From this view, you can see how complete the orders in a folder are, how complete a specific order is, and whether it is early, on time, or late. The system distinguishes critical routings and operations. The system does not display work orders that are complete or cancelled.

The nested structure of the Work Order Centric view also displays multilevel pegging between work orders, net new production recommendations, and purchase order recommendations, as shown in the following screen shot:
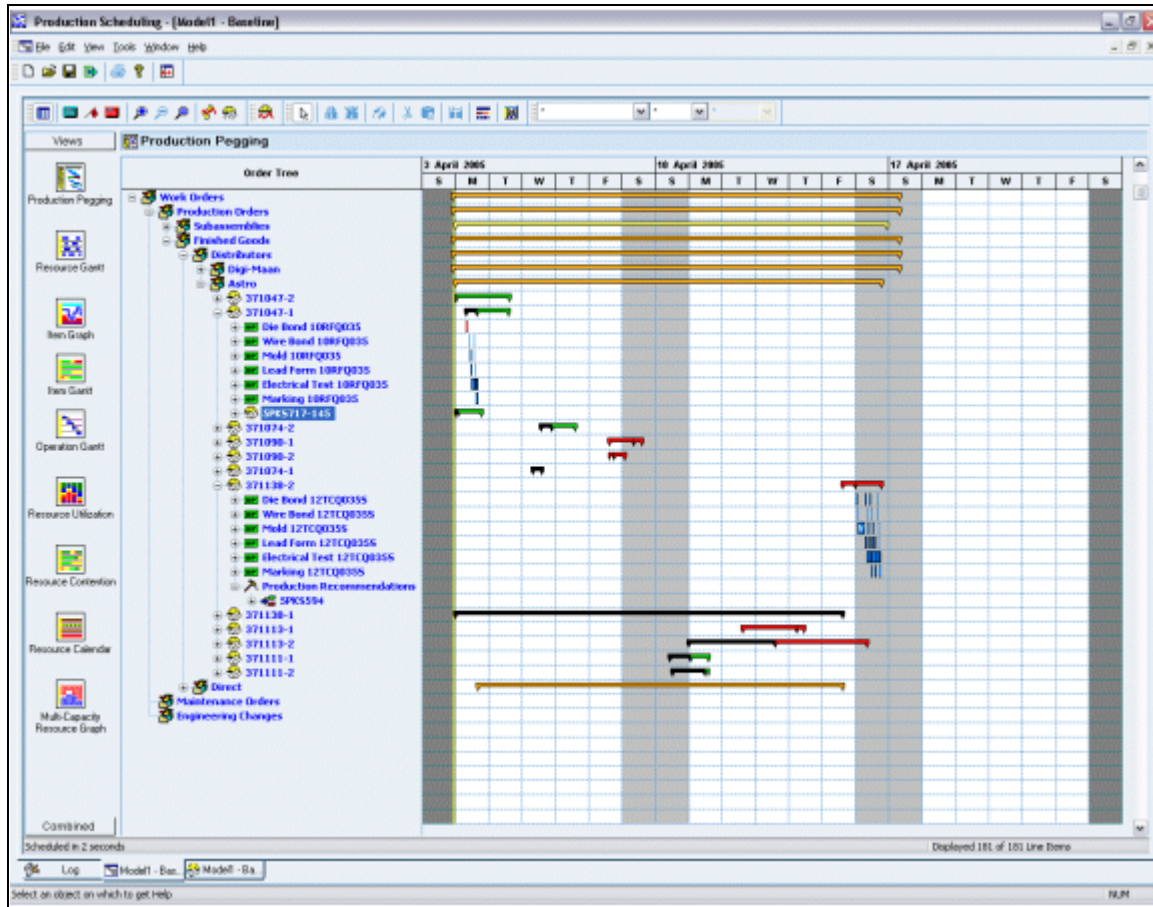
Production Pegging - Work Order centric view

Production recommendations display Production Scheduling suggested production requirements when your work order quantities are insufficient to fulfill demand or upstream/downstream production requirements.

The Production Pegging - Work Order Centric view has the same folder structure as the Work Order Editor and allows you to drill down into the Work Order Editor from the pegging view.

To open the Work Order Editor from the Production Pegging - Work Order Centric view:

1. Select a work order in the work order tree.

2. Click the Open Work Order Editor icon in the menu bar.

The Work Order Editor opens the work order you selected and highlights the work order number in the Work Order tree, as shown below:

Work Order Editor highlighting a selected work order

The Production Pegging - Work Order centric toolbar provides a number of buttons that are available only in Production Pegging. From this toolbar, you can filter work orders according to their criticality to assist you in analyzing the schedule's handling of orders. You can filter work orders to list only those of a specific type, punctuality, and magnitude of punctuality. Production Pegging retains the last filtering settings and uses them for future analysis.

Orders can be filtered to display only those that meet the specific type of criteria that you specify. This information can help you locate any issues and resolve them by modifying the resources or materials involved.

The following criteria are used to filter work orders:

Filter	Name	Description
Primary	Work order filter	Enables the work orders to be filtered by the type of work order. Options are: Production, Maintenance, or Engineering.
Secondary	Order Punctuality filter	Enables the orders to be filtered by punctuality. Options include all, early, on time, and late.
Tertiary	Order Magnitude Filter	Enables the user to specify the magnitude of punctuality when filtering. By default, these options include all, >0.5 days, >1 day, >2 days, >3 days and >5 days. Select Customize to set your own time filters.

Note. The number of late orders is shown in the status bar at the bottom of the screen when you use the filtering in the view.

Operations

For each demand line item and routing, the related operations are displayed. The color of each operation is determined by whether it is a critical bottleneck operation. The start time of critical operations is the same as their earliest start time. No flexibility exists for the operation to start any earlier than its current start time. The start time for noncritical operations is later than their earliest start time.

Color	Significance
Blue box	Noncritical operation.
Red box	Critical operation.

A pop-up box displays the following information:

Operation name	Name of the operation.
Start time	Start time for the operation.
End time	End time for the operation.
Duration	Length of time between the start time and the end time.

Resources

The resources for each operation can be viewed in the Production Pegging. These details are similar to those displayed in the Resource Gantt view. The Time Grid displays each instance that the resource is used for an operation.

Element	Appearance
Single capacity resource	Blue with operation name.
Multicapacity resource	Multiple colors, depending on the percentage of capacity being used. Ranges from dark green for low usage to red for full capacity use.
Changeover	Thin black bar.

Supply Events

At times, you might need to see the demands that are dependent on specific supply events. From Production Pegging, you can peg up from a supply event to all of the associated demands.

Supply events are displayed as diamonds in the Time Grid. The color green indicates any supply event.

Pop-up windows are also available for supply events. The following fields are included:

Supply Number	Supply event order and line number.
Item	Inventory item number on the supply order.
Quantity	Quantity of item that is requested.
Requested Date	Date when the item is expected to be fulfilled.

Partially Filled Demand

If an order is filled from production and either starting inventory or a supply event, the starting inventory is displayed in the Production Pegging view by a green diamond. The green diamond at the start of the horizon indicates that some of the order is taken from inventory. The production below the green diamond row indicates that the remaining quantity is from manufacturing activities. Bubble help provides additional details for each event such as pegged quantity, date, and so on.

Purchase Recommendations

Purchase recommendation are recommended purchase activities for produced items. They appear as a green diamond with a tail that represents the lead time from the supplier. Placing your mouse pointer over the purchase recommendation provides this information:
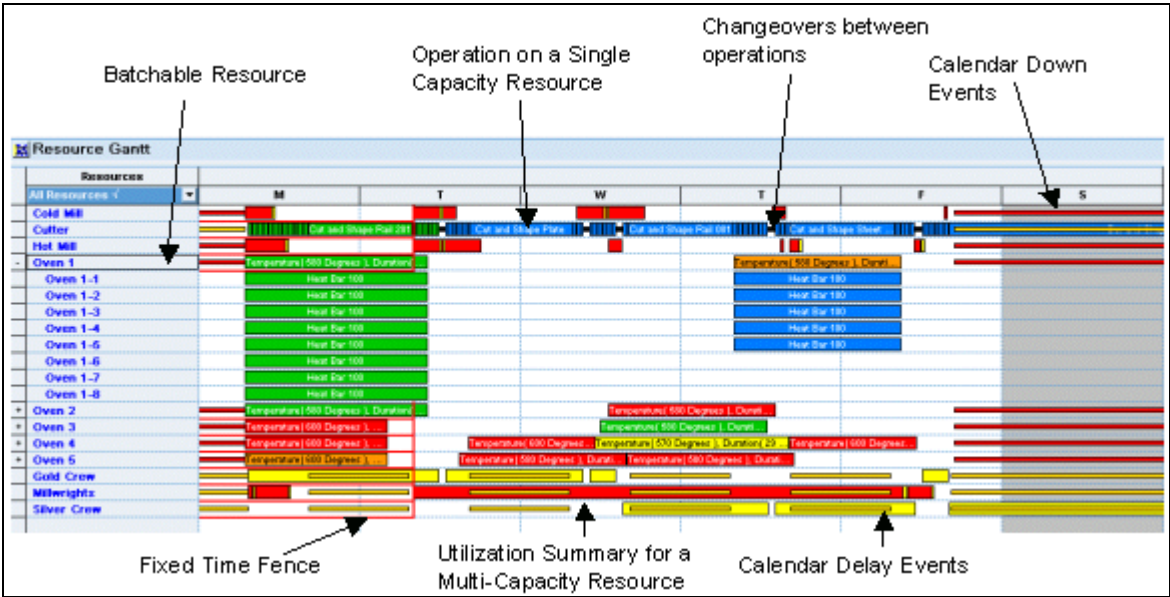
Supplier	Displays the selected supplier.
Item	Displays the purchased item.
Pegged Quantity	Displays the quantity that is required for the order.
PO Recommendation Quantity	Displays the amount ordered that may also supply other orders.
Required Date	Displays when the purchase order is needed for production.
Lead Time	Displays the lead time from the supplier.
Estimated Ship Date	Displays the date that the purchase order is estimated to ship from the supplier.

The Resource Gantt View

The Resource Gantt is a resource-centric view of the schedule. It enables the planner to see exactly how machines, crews, and tools are scheduled across the horizon. Operations are displayed with blue rectangles to indicate exactly where they start and end. You can zoom into any point in the horizon by selecting a range in the upper portion of the horizon bar and then select the Zoom In button on the toolbar.

For multicapacity resources, color is used to represent the percentage of capacity currently being used. The Set Colors button defines these colors. For example, if you have not changed the default color setting, dark green represents low (0 to 10 percent) capacity used, while red represents high capacity (90 to 100 percent).

This page shot illustrates the chart elements that are described in the following table:

Resource Gantt chart elements

This table describes the chart elements:

Element	Appearance
Batchable Resource	A summary level row which displays information on the planned batch and a little *plus* arrow beside it that the user can expand or collapse. This provides visibility into the spot usage of the batch. For example, how many operations, which operations, and so on.
Changeover	Thin black bar.
Delay	Thin yellow bar.
Down time	Thin red bar.
Fixed Production Timefence	Thin red line from the start of the horizon on the resource which extends to the end of the fixed production timefence.
Multi-capacity resource	Multiple colors, depending on the percentage of capacity being used. Ranges from dark green for low usage to red for full capacity use.
Single capacity resource	Blue with operation name.
Campaign Run Optimization Cycle Boundaries	Thin brown bar. This represents manufacturing cycle boundaries when using Campaign Run Optimization. This provides for improved visibility to changeover patterns in and across manufacturing cycle boundaries.

You can view all resources in the model at once, or display resources of a particular type. Resource usage is displayed for the entire schedule horizon, and can be zoomed to display a particular time period. You can view any resource group in the Resource Gantt, Resource Utilization, and Resource Contention views. The attributes of the resource group must be appropriately set for the group to be accessible in the Resource Gantt. If the group does not appear in the list, verify that the properties of the group are set to display in the Resource Gantt. If you add a group after the initial solve, you need to close and reopen the schedule, and then open the relevant view.

The default Resource Gantt sort order is alphanumeric by resource type. To more easily locate potential resource conflicts, you can sort the Resource Gantt view based on a resource's utilization as a percentage of its total capacity. When you sort by utilization, resources with the highest utilization appear at the top of the Resource Gantt. You can also change the view back to the default sort order.

Note. If a resource is campaign run optimized, its name appears in green.

You can use a variety of manual scheduling options in the Resource Gantt to fine-tune your schedule.

If you right-click on a resource, you have several options:

- You can drill down to the Resource Properties.

- You can relax constraints, which allows capacity and calendar violations during the solve process.

- You can reactivate the calendar and capacity constraints on the resources by performing a right- click and selecting Active Constraints.

- You can select Enforce Horizon Start, which applies to both the cold solve and the repair solve.

  This flag automatically reconciles any shop floor backlog, when scheduling, by ensuring all operations start at the beginning of the horizon, even if they have firm dates that are prior to the start of the horizon. The Enforce Horizon Start option enables you to move operations running in the past, as a result of shop floor backlog, to the start of the horizon. Selecting the Enforce Horizon Start option and performing either a repair solve or a cold solve pushes forward all operations on that resource to the start of the horizon, maintains operation sequence, and respects upstream and downstream precedence constraints.

See "Defining Resource Availability," Creating Production Scheduling Models, Defining Resources, *Production Scheduling*

See "Using the Enforce Horizon Start Option," Manual Scheduling, Understanding Solved Schedules, Operation Adjustments, *Production Scheduling*

Batchable Resources

When evaluating a batchable resource in the Resource Gantt, color indicates utilization at the batch level and bubble help indicates the attributes of the batch at the resource level. You can double-click a resource row to expand or collapse visualization of individual spot usage. The toolbar provides an option to highlight certain batches based on the resource attribute type and attribute value. The selected batches are highlighted in orange.

See "Batchable Resources," Creating Production Scheduling Models, Defining Resources, Understanding Resources, *Production Scheduling*

The Item Graph View

The Item Graph presents schedule information on item levels in your supply chain using a line graph. This view is a line graph that displays current levels for each item in your model. On the Item Graph, you can select any item and then view levels for that item by moving your cursor to any point on the line.

Using the item graph, you can view inventory levels for a specific item at a specific time and access detailed property information for the item. You can customize the color scheme used in the Item Graph by selecting the color pallet icon in the toolbar.

Note. If you select Relaxing Minimum, Relaxing Maximum, or both in Item Properties, you are giving the Production Scheduling solver permission to temporarily violate item minimum or item maximum levels, respectively, in order to produce a feasible solution.

The color in which an item is displayed depends on whether you have relaxed the minimum constraints, maximum constraints, or both constraints in Item Properties. The colors are outlined in the following table:

Option	*Maximum Constraints*	*Minimum Constraints*
Relaxed	Default setting. Items are displayed in dark blue. If minimum constraints for an item are also set to relaxed, items are displayed in orange	Items are displayed in the color you chose for item minimum levels. If maximum constraints for an item are also set to relaxed, items are displayed in orange
Unrelaxed	Items are displayed in the same color as the color you chose for item maximum levels.	Items are displayed in dark blue.

The Item Gantt View

The Item Gantt is a bar chart that displays colored bars representing item levels. The Item Gantt displays a horizontal bar chart that represents items and storage spaces, and how they relate to the defined maximum and minimum levels. It can be filtered for specific resource groups. As with other schedule views, you can view current information at different levels of detail by using the zoom tools. Color-coding enables you to see the current status and problem areas immediately.

From the Item Gantt, you can view the inventory level of each item in relation to its minimum and maximum levels, and also access an item's property pages.

You can view any resource group in the Item Gantt. An item can be in more than one resource group. You can create as many resource groups as necessary to simplify your schedule analysis. To be accessible for the Item Gantt schedule view, the resource group's attributes must be appropriately set.

To view items that could inhibit the smooth implementation of your schedule, you can view the Item Gantt sorted by degree of violation in descending order. Sorting by degree of violation puts the most problematic items at the top of the view for further review and fine-tuning. Items are sorted in the following order:

1. Items with at least one red (inventory that is below minimum levels) segment.

2. Items with at least one violet segment, that is, inventory that is greater than the maximum level, and no red segments.

3. Items with only orange (inventory that is within 10 percent of the maximum or minimum levels) and green segments.

4. Items with only green segments.

Note. To get a better view of an item's inventory levels, create a combined Item Gantt and Item Graph view.

The Item Gantt uses colors to represent item levels in relation to their minimum and maximum levels. You can click the Set Colors button to customize the colors that are used in the Item Gantt and to set the thresholds for each item. The following table is a summary of the colors that are used in the Item Gantt:

Color	*Definition*
Green	The item level is less than 10 percent of the difference between the minimum and maximum levels.
Yellow	The item level is greater than 10 percent of the difference between the minimum and maximum levels.
Red	The item level is below the minimum level.
Violet	The item level is greater than the maximum level.

The Operation Gantt View

The Operation Gantt lists all of the scheduled operations. It is typically used in conjunction with the Resource Gantt in the form of a combined view as the contents of the Operation Gantt become context sensitive to what is selected in the Resource Gantt. For example, when selecting a resource in the Resource Gantt, only operations running on the selected resource are displayed in the Operation Gantt: top to bottom, left to right, by start date. Each operation instance in the Operation Gantt indicates when the operation is run, and its duration, including any setup and delay times. You can expand the list of operations to display every operation instance, or you can create operation groups and filter operations by these groups. Pegging information is also available from this view. You can use the Operation Gantt to review operation properties. Right-click a displayed operation.

You can use the manual scheduling tools in the Operation Gantt to make detailed adjustments to your schedule. The Remove Idle Time command enables you to remove idle time between two or more selected operations or all idle time that precedes a selected operation. You can also manually resequence operations using the cut and paste commands.

You can also open the applicable operation or routing diagram from the right-click menu.

Note. The cut, paste, and Remove Idle Time commands are enabled only for multiple instances of the same operation.

The Resource Utilization View

The Resource Utilization view provides you with precise information about all of the defined resources that are used in the schedule horizon. Information such as idle time, calendar down time, delay time, changeover time, and run time is displayed in this view. Resource utilization appears as total number of hours utilized and as a percentage. By choosing from drop-down menus, you can display each resource and its utilization by shift, day, week, or month time buckets. If you select a column in the utilization view, summary information for that time period is displayed in the left side of the screen for that specific time bucket.

The following table describes each item in the legend:

Legend Entry	Description
Idle	The percentage of time and total hours when the resource is available but not utilized.
Delay	The percentage of time and total hours when the resource is not used because of a delay.
Down	The percentage of time and total hours when the resource is not used because of down time.
ChOvr	The percentage of time and total hours that the resource is not used because of a changeover.
Run	The percentage of time and total hours when the resource is being run.
Start Time	The date and time of the start of the utilization time period.
End Time	The date and time of the end of the utilization time period.

Note. If your schedule contains changeovers, delays, or down time, utilization statistics appear in the legend and the corresponding color-coding appears in the columns. If your schedule does not contain changeovers, delays, or down time, these events appear as zero percent in the legend.

The Resource Contention View

The Resource Contention view captures the contention encountered for each resource throughout the solve process. This information is useful to find the dominating item or resource bottlenecks during the solve process. The view also serves as a histogram of resource contention if you navigate to it after a solve.

By default, the Resource Contention view displays the contention events that are encountered during the solve using the time bucket.

Schedule Horizon	Default Time Bucket
Greater than 3 months	Month

Schedule Horizon	Default Time Bucket
Between 2 weeks and 3 months	Weeks
Less than 2 weeks	Days

You can change the time bucket used by the Resource Contention view to calculate the contention. Time bucket options include month, week, day, shift, and hour. The time bucket setting is independent of the view that is currently displayed using the zoom buttons.

For example, if you change your time bucket to the day option, then the Resource Contention view evaluates and compares the contention during each day, and displays its findings using the color scale. The colors are representative of the contention for each resource over each day. The smaller the time bucket, the more differentiated the results are, due to the smaller pool of data.

Color is used to represent contention intensity. The Set Colors button defines these colors. There are 10 ranges of colors in 10 percent increments.

Using the Set Colors command, you can set new color levels by adding new sliders or moving existing sliders in either direction. This enables you to assign a color to represent different percentages of contention. The colors that you set in the color bar are displayed in schedule views, so that you can always see a consistent color representation of the schedule.

The Multi-Capacity Resource Graph View

The Multi-Capacity Resource Graph displays a resource's total available production capacity and how much of that capacity is used over time. You can select any resource and then view levels for that resource by moving your pointer to any point on the graph.

The area in red, at the top of the graph, represents the maximum capacity of the resource. The red shaded area, at the bottom of the graph, represents zero utilization point. You can customize the colors in the view using the icon on the toolbar.

Combined Views

To get more detailed information about a generated schedule than from a single schedule view, you can select a combination view. With a combination view, two or more views are synchronized so that when an item is selected in the main view, the other view displays the related information only about that item. For instance, while reviewing the Resource Gantt, you can select a resource and simultaneously view its utilization, item usage, or operation details during a specific time frame. This powerful feature enables you to quickly analyze key elements of your schedule and determine whether any adjustments need to be made. The following combination views are available:

Resource and Operation Gantt   You can see operations that run on selected resources in the Resource Gantt.

Resource Gantt and Utilization   You can see the utilization of selected resources in the Resource Gantt.

Item Gantt and Graph   You can see the inventory levels of selected items in the Item Gantt.

Resource and Item Graph	You can see the inventory level of an item that is produced or consumed by a single operation.
Resource Gantt and Multi-Capacity Resource Graph	You can see the utilization of selected multi-capacity resources.
Resource Gantt, Operation Gantt and Item Graph	You can see operations that run on selected resources in the Resource Gantt, and you can see the inventory level of an item that is produced or consumed by a single operation by selecting an operation in the Resource or Operation Gantt.

Additional combination views can be created and added to the combined group, if necessary. Each view can have the name and button of your choice. To organize your views, you can also set up group views that contain specific schedule views.

While in this view, you can use the scroll bars and zoom options to move to different areas in your view. As long as the synchronized scroll bars feature is active, both views stay synchronized as you move around in the view.

Filtering Work Orders

This section discusses how to:

- Use the Work Order Advanced Filter

- Understand filtered views

- Add a new filter

- Delete a filter

- View and Modify filter properties

- Save work order filters

- Load work order filters

- Assign a filter to a view

- Remove a filter from a view

Using the Work Order Advanced Filter

The Work Order Advance Filter allows you to further "slice and dice" data in the various Gantt, Utilization, and Pegging views in a multi-dimensional fashion and filter the content of these views. This feature allows the planner to hone in on specific areas of the schedule only. For example, you can create a filter to "show me all work orders in from October 25 – November 5th running on Resource (or Group) which have a Firm Start status". The results are previewed in the bottom of the filtering view and you can drill-down to the Work Order Editor directly from here as well if desired. This filter is especially useful for Maintenance Work Orders scheduled from cMRO as filtering can be done on the Work Breakdown Structure.

The Work Order Advanced Filter window is used to view, create, and modify filters for various Production Scheduling views. Once the filter is defined, it can be applied as need to the views.

For example, you may wish to filter data such as:

- All Work Orders for a Visit.

- All Work Orders for a stage.

- All Work Orders for a specific department.

- All Work Orders for a date range (beyond standard zooming).

- All Work Orders by statuses and priority.

- All Work Order assignments by employee or crew.

- All Work Orders by cMRO parameters.

If no filters current exist in your model, click the Toggle Work Order Filter button on the toolbar.

If filters have been defined in your model, select refer to Adding a New Filter or Viewing or Modifying Filter Properties for accessing the Work Order Advance Filter window.

From the Work Order Advanced Filter window, choose the criteria for the filter using the options displayed within the window. The more filters selected, the narrower the filtered results will be.

The Work Order Advanced Filter fields are as follows:

Filter Name	Enter the name you want to assign to the filter.
Created in scenario	A display only field that shows the scenario used to originally define the filter.
Work Order Number	To filter for a specific work order, select the work order from the list. All work orders within the model are displayed in the list.
Work Order Description	Enter text to perform a query based on text within the work order Description field. You can use the "%" wildcard symbol for contextual text filtering. For example, all work orders which contain the word avionics would be represented as %avionics%.
Work Order Status	To create a filter based on work order status, select the appropriate option from the list. This field should be dynamic based on the populated work order status values in the model.
Produced Item	This is a list of items in the model and should include the cMRO maintenance items.
Order Class	Use to specify Production Scheduling order class value (Hot, Committed, Uncommitted).
Priority	Use to specify work order prior. Valid values are $1 - 5$.

Type   Use to specify work order type. Valid values are Production, Maintenance, and Engineering.

For the Work Order Quantity Ranges region, the fields are as follows:

Quantity Required   Specify the quantity required range.

Quantity Complete   Specify the quantity complete range for the filter.

Quantity Remaining   Specify the quantity remaining range for the filter.

For the Work Order Date Ranges region, the fields are as follows:

Work Order Request Date   Use to specify request date range.

Work Order Start Date   Use the specify work order start date range.

Work Order Completion Date   Use the specify work order completion date range.

Use the following fields to define a filter based on manufacturing data for the work orders. For the Manufacturing region, the fields are as follows:

Routing   Use the list to define a routing for the filter. Work orders assigned to the routing will appear based on filter criteria.

Original Routing   In the event a work order was offloaded to an alternate routing, select the original routing option assigned to the work order.

Contains Operation   Select the all operations in the model. The operation list should be filtered by the routings chosen (if populated).

Contains Operation Status   Use the list to assign an operation status to the filter. Valid options are Open, Active, and Closed.

Uses Resource / Group   Lists all resources and resource groups in the model. If routings and/or operations are selected above, the values in this list should be limited to the resources which are consumes in those routings and/or operations.

Consumes Material / Group   Lists all items and item groups in the model. If routings and/or operations are selected above, the values in this list should be limited to the resources which are consumes in those routings and/or operations.

Contains Firm Status   Use to perform a query based on work order that have an assigned firm status. Valid values are Firm Start, Firm End, and Firm Range.

Firm Start Date   Use to define a firm start data range for your filter.

Firm End	Use to define a firm end data range for your filter.

Use the following fields to define a filter based on cMRO data for the maintenance work orders. For the CMRO region, the fields are as follows:

Visit	Use to perform query on cMRO work orders that use this field. List displays all the visits assigned to work orders in the model.
Stage (Attribute 1)	Use to perform query on cMRO work orders that use this field. List displays all the stages assigned to work orders in the model.
Process (Attribute 2)	Use to perform query on cMRO work orders that use this field. List displays all the processes assigned to work orders in the model.
System (Attribute 3)	Use to perform query on cMRO work orders that use this field. List displays all the systems assigned to work orders in the model.
Access Panel (Attribute 4)	Use to perform query on cMRO work orders that use this field. List displays all the access panels assigned to work orders in the model.
Zone (Attribute 5)	Use to perform query on cMRO work orders that use this field. List displays all the zones assigned to work orders in the model.
Subzone (Attribute 6)	Use to perform query on cMRO work orders that use this field. List displays all the subzones assigned to work orders in the model.
Department (Attribute 7)	Use to perform query on cMRO work orders that use this field. List displays all the departments assigned to work orders in the model.
Unit Name (Attribute 8)	Use to perform query on cMRO work orders that use this field. List displays all the unit names assigned to work orders in the model.
Non Routine (Attribute 9)	Use to perform query on cMRO work orders that use this field. Valid options are Yes, No, or Null. This option is used indicate whether the work order is of type non routine or not. If set to "null", all work order appear. If set to Yes or No, then filter performs query based on status as selected.

All work orders that meet the defined filter criteria appear in the Search results region at the bottom of the Work Order Advanced Filter window. You can view a work order in the Work Order Editor by double-clicking on the work order row.

Understanding Filtered Views

Once a filter is defined, it can be assigned to the following views:

- Production Pegging – Only work orders which meet the criteria specified should be displayed. Empty folders should not be displayed. In the event a work order has multiple operations and resource level filtering is specified, only the work order operations which meet those criteria are outlined.

- Resource Gantt – Only operations belonging to the work orders in the filter are displayed. This includes the utilization shading for multi-capacity resources.

- Operation Gantt - Only operations belonging to the work orders in the filter are displayed. The Operation Gantt should also be filtered to only operations running on this resource.

- Resource Utilization – Utilization should be calculated based on filtered work orders. This enables users to understand utilization contribution.

- Multi-Capacity Resource Utilization - Utilization should be calculated based on filtered work orders. This enables users to understand utilization contribution.

Adding a New Filter

To create and save filters to your scenario using the Work Order Advanced Filter window:

1. Select New Filter from the list option displayed to the right of the Toggle Work Order Filter button on the toolbar.

2. Define the filter as needed from the Work Order Advance Filter window.

3. If you want to save the filter, click OK. To exit the Work Order Advance Filter window without saving the filter, click Cancel.

Deleting a Filter

To delete a filter:

1. Select Work Order Filters from the list option displayed to the right of the Toggle Work Order Filter button on the toolbar.

2. Select the filter you want to remove from the Work Order Filters dialog and click Delete.

3. Click Close.

You can also delete a filter by selecting the filter from the from the list option displayed to the right of the Toggle Work Order Filter button on the toolbar. A check mark appears when the filter is selected. Delete the selected filter by selecting Delete Filter from the Work Order Filters list option. The filter is removed from the scenario.

Viewing and Modifying Filter Properties

To view and modify filter properties:

1. Select Work Order Filters from the list option displayed to the right of the Toggle Work Order Filter button on the toolbar.

2. Select the filter you want to view and click Properties. The Work Order Advanced Filter window appears and displays the filter criteria.

3. If you want to make changes to the filter, adjust the fields as necessary and click OK to save your changes.

Saving Work Order Filters

Once you define a filter, you can save your filter definitions to a Work Order Filter (.wof) file. This file can then be used to load and apply the filters to other Production Scheduling models and scenarios as needed.

To save work order filters:

1. Select Work Order Filters from the list option displayed to the right of the Toggle Work Order Filter button on the toolbar.

2. Click Save As. The Save Work Order Filters As dialog appears.

3. Enter the file name you want to assign in the File name field.

4. Navigate to the location where you want to save. By default, the file is saved to the same location as the model used to define the filter.

5. Click Save.

Loading Work Order Filters

Production Scheduling filters defined in Work Order Filter (.wof) files can be loaded and used in your model. Once loaded into your model, you can adjust the filter as needed based on your needs.

To load work order filters:

1. Select Work Order Filters from the list option displayed to the right of the Toggle Work Order Filter button on the toolbar.

2. Click Load. The Load Work Order Filter dialog appears.

3. Navigate to and select the appropriate file.

4. Click Open. All filters defined in the file appear in the Work Order Filters dialog.

Assigning a Filter to a View

Once a filter is assigned to your model, you can apply it when needed.

1. Open the appropriate view.

2. Select the filter you want to apply from the list options displayed to the right of the Toggle Work Order Filter button on the toolbar. A check mark appears next to the selected filter.

   The Toggle Work Order Filter button appears highlighted when a filter is applied to a view and the selected filter name appears to the right side of the view title bar.

Removing a Filter from a View

Click the Toggle Work Filter button on the toolbar. You can also remove a filter by deselecting the filter from the list option displayed to the right of the Toggle Work Order Filter button on the toolbar. The view is refreshed to display all data.

Analyzing Scheduled Resources

This section discusses how to:

• Relax resource constraints.

• View pegging resources.

• Sort resources.

• View resource utilization.

• View resource contention.

• View alternate resources.

Relaxing Resource Constraints

Resource constraints can be relaxed and reactivated when viewing the Resource Gantt, Resource Utilization, and Resource Calendar.

To relax resource constraints:

1. Access the appropriate schedule view.

2. Right-click the resource and select Relax Constraints.

   The relaxed resources appear in orange in the schedule views.

3. Click Solve to solve the schedule.

To select more than one resource at a time, hold down the Ctrl key to select individual resources, or press Shift to select a block of resources.

To reactivate resource constraints right-click the resource and select Activate Constraints.

Viewing Pegged Resources

To view pegged resources:

1. Access the Resource Gantt.

2. Select one of the following options:

   • *Peg Upstream,* which displays upstream operations.

   • *Peg Downstream,* which displays downstream operations.

   > **Note.** To display both upstream and downstream operations, select both options.

3. Click the first operation for which you want to view pegging information.

   The Resource Gantt highlights the upstream operations in brown and the downstream operations in purple.

4. Repeat the previous step for other operations to view pegging information.

5. Click the pegging buttons again to turn off this option.

Sorting Resources

To sort resources:

1. Access the Resource Gantt.

2. On the schedule view toolbar, click Zoom In to increase the size of the columns until dates are displayed across the top.

3. On the sort order drop-down list box, select one of the following options:

   • *Default Sort,* which sorts resources alphanumerically.

   • *Utilization Sort,* which reorders the resources in the Resource Gantt in order of their utilization. Resources are sorted by utilization from highest to lowest.

   • *Peg Sort,* which reorders the operations in the Resource Gantt according to the routing order and their associated resources. You can turn off this feature when it is not required. Doing so returns the resource order to its original sequence.

Viewing Resource Utilization

To view resource utilization:

1. Access the Resource Utilization view.

2. Click Zoom In, Zoom Out, or Zoom to 100 percent to view the schedule at the required level of detail.

3. Select a resource from the Resources list box.

   The vertical bars in the graph represent the utilization level for that resource.

4. Select one of the following time periods from the Bucket Set available options: *Shift, Daily, Weekly, Monthly.*

5. In the graph, click a time period to display utilization statistics for that time period.

Viewing Resource Contention

To view resource contention:

1. Access the Resource Contention view.

2. Click Zoom In, Zoom Out, or Zoom to 100 percent to view the schedule at the required level of detail.

3. Select a resource from the Resources available options.

   The horizontal bars in the graph represent the contention level for that resource.

   Note. You should refer to the Resource Contention view after the schedule has been solved. Watching the Resource Contention view during the solve slows down the solve because of the necessary graphical updates.

Viewing Alternate Resources

To view alternate resources:

1. Access the Production Pegging view.

2. Select a resource from the Demand Tree.

3. Right-click and select Show Resource Alternates.

   Alternates for the specific resource are displayed in the Production Pegging view at the resource level when you expand an operation. You can also manually offload to an alternate using cut and paste, or drag and drop, and running a repair solve.

Analyzing Scheduled Items

This section discusses how to:

- View item graph data.

- View item properties.

- View item groups.

- View item production and consumption.

- Sort by degree of violation.

Viewing Item Graph Data

To view item graph data:

1. Access the Item Graph.

2. Select one of the following options from the Items list box:

    •   All Items

    •   An item group

    •   An individual item

3. On the Items list, click an item.

    The chosen item is selected, and the level for that item displays as a line in the graph.

4. Move your cursor over any point on the chart line.

    A pop-up box displays item description, current quantities, and level start and end dates.

5. Click the zoom buttons to display more or less detail.

6. On the toolbar, click the item graph buttons as follows:

    •   Click the Scale buttons to change the scaling on the y axis.

    •   Click the Show Spikes button to display overlapping supply and demand events that occur at the same time.

Viewing Item Properties

To view item properties:

1. Access the Item Graph or Item Gantt.

2. Use the Items list box to view a specific item or display a group of items.

3. Right-click an item and select Properties.

Viewing Item Groups

To view item groups:

1. Access the Item Gantt.

2. Click All Items.

3. Select the specific item group that you want displayed in the Item Gantt.

    Note. If the group does not appear in the list, verify that the properties of the group are set to display in the Item Gantt. If you add a group after the initial solve, close and reopen the schedule, and then select the Item Gantt.

Viewing Item Production and Consumption

To view item production and consumption:

1. Access a Resource and Operation Gantt view combined with the Item Gantt and Graph view.

2. Click the operation for which item data is displayed.

3. Click the Produced Items or Consumed Items button from the toolbar.

   The Item Gantt (or Graph) view is filtered for the option that is selected.

Sorting by Degree of Violation

To sort the degree of violation:

1. Access the Item Gantt.

2. Click Sort Items.

   The items are sorted in descending order with the most contentious items at the top of the view.

3. To return the view back to its original sort order, click Sort Items again.

   The view returns to its original appearance.

Analyzing Scheduled Operations

This section discusses how to:

• View an operation quick summary.

• View operation groups.

• View pegged operations.

• Modifying firm status and offloading setting.

Viewing an Operation Quick Summary

To view an operation quick summary:

1. Access the Operation Gantt.

2. Hold your cursor over any operation.

   Information about the operation, including its name, start time, and duration, appears in a pop-up window.

Note. This procedure can be used to view quick summary information in the Production Pegging and Resource Gantt.

Viewing Operation Groups

To view operation groups:

1. Access the Operation Gantt.

2. On the Operation Gantt, select the operation group that you want to view from the Operations available options.

3. Obtain schedule information as described in the following steps.

4. Click the Alerts button to view any alert messages.

5. Click the zoom buttons to view different levels of detail.

6. Right-click any operation to view and make changes to properties.

   Note. If you rename the group or change properties, you need to close and reopen the schedule, and then select the Operation Gantt.

   If the group does not appear in the list, verify that the group properties are set to display in the Operation Gantt.

Viewing Pegged Operations

To view pegged operations:

1. Access the Operation Gantt.

2. Select one of the following options:

   • To display pegged operations upstream of the selected operation, click Peg Upstream.

   • To display pegged operations downstream of the selected operation, click Peg Downstream.

   • To display all pegged operations, click the Peg Upstream or Peg Downstream icons on the toolbar.

3. Select the first operation for which you want to display pegging information.

   The Operation Gantt view highlights the upstream operations in brown and the downstream operations in purple.

Modifying Firm Status and Offloading Setting

To adjusting firm status and offloading:

1. Access the Operation Gantt.

2. Select an operation.

3. Right-click and select the appropriate firm status. The following options are available:

Firm Start	Select this option to use a firm start date.
Firm End	Select this option to use a firm end date.
Firm Range	Select this option to use a firm date range (start date and end date).
Unfirm	Select this option to remove any firm settings.

Firmed operations appear green in the view. To specify a specific firm dates, go to the Routing tab in the Work Order Editor and modify the Start Date and End Date fields as needed.

4. To adjust the offloading setting, select the appropriate offload option. The following options are available:

Allow Offload	Select this option to allow operation to be offloaded to an alternate resource.
Disallow Offload	Select this option to prevent an operation from using an alternate resource.

5. Perform a solve to update your schedule.

You can also elect to manually populate the firm dates from the Resource tab in the Work Order Editor if desired. If you select Firm Range from the Gantt view, Firm End the appropriate values and dates would show up in the Work Order Editor, the main difference being which date(s) are populated. (Firm Start populates Start Date of the Work Order Operation, Firm End populates the End Date and Firm Range populates both start and end dates).

You can also set the offloading from the Bill of Resource tab in the Work Order Editor. Select No from the Allow Offloading list to prevent offloading from occurring.

Using Combined Views

This section discusses how to:

• Add a combined view

• View combined views

Adding a Combination View

To add a combination view:

1. Access the View Wizard window.

2. In the List of available subviews, select a schedule view that you want to appear in your combination view.

3. Click Next to move the subview to the Selected subviews area.

4. Repeat steps 1 and 2 for each additional view.

5. If necessary, click the Set Order buttons to delete or reorder views.

    When the new view is created, the views are displayed in the same order that they are displayed in the Selected subviews list. A full description of each button is described in the following table:

Button or Icon	*Name*	*Description*
Button	Delete view	If a view has been moved into the Selected subviews list by mistake, select the view, and then click this button to remove it from the list.
Icon	Move up view	To move a view up in the Selected subviews list, select the view, and then click this icon to move it up in the list.
Icon	Move down view	To move a view down in the Selected subviews list, select the view, and then click this icon to move it down in the list.

6. Click Next, enter a view name, and click Next again.

7. On the View Wizard, select an button for the group, and click End.

    The button that you select and the name is added to the bottom of the current view group on the screen.

Viewing Combined Views

Use this page to view combined views:
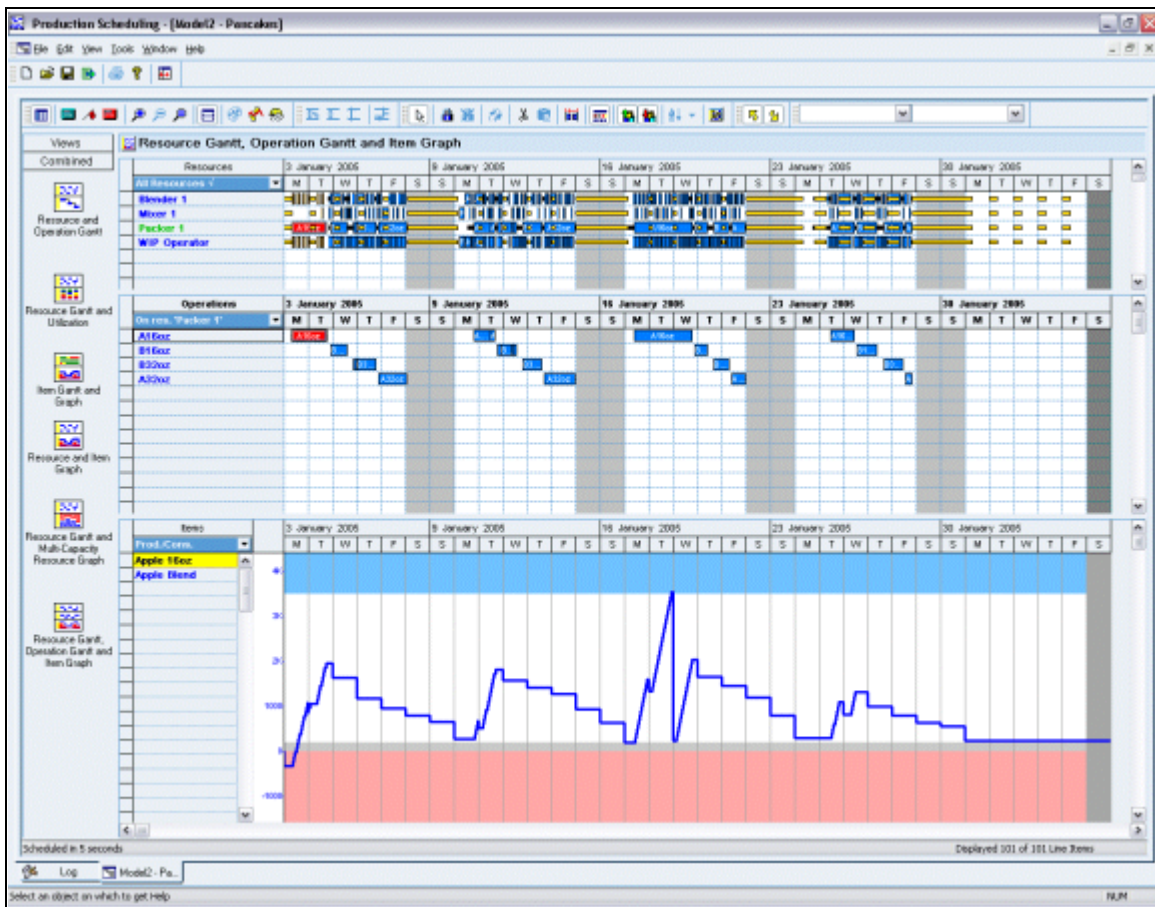
Combined view

To view combined views:

1. Access the List Bar.

2. Click the button for the combined view that you want.

3. Click on an element in the main view for which you want more details.

In a Resource Gantt, you can click any resource. The detail view might display how this resource is being used or the operations to which it is associated, depending on your choice of detail view.

In the Resource or Operation Gantt views, you can right-click any operation and display either the produced or consumed items, or all of the items that are produced and consumed by the operation.

In the Item Gantt, you can click any item. The detail view (the Item Graph) displays further information about that item's stock levels.

Working with Alerts

This section provides an overview on alerts and discusses how to:

•   Access alert messages.

- Navigate to production pegging alerts.

- Filter alert messages

Understanding Alerts

Alerts are available from either a solved or unsolved schedule. Alerts display any violations of soft constraints. Alerts can be accessed from the model workspace in a tree format. When you click an alert that is displayed on the tree, the Alerts grid opens in the bottom of the applicable view in which you are working. The Alerts grid contains tabs that break out alerts by category including Demand alerts, Inventory alerts, Work Order alerts, and Resource alerts. Within each of these categories are various types of alerts. The alerts grid displays information about each violation so that you can resolve the problem that is identified.

The Root Cause column of the alerts grid provides information to guide you to any problem operations. For instance, a message such as "Purchased Part Constraint for Deliver 9011" indicates a possible shortage of purchased parts on hand for the operation called Deliver 9011. To further expand on the operation itself, you can view the operation properties directly from the alerts grid. The Properties dialog box displays the time when the operation is scheduled to run and its associated resources. Collectively, this feedback can help to identify the root of the problem.

You can navigate from a demand or work order alert, in the alerts grid, to view the violation in the Production Pegging view. When you navigate to the Production Pegging view, the applicable demand or work order is displayed and the operation in question is highlighted.

The Alerts Tree

In the Model Workspace, the alert tree is broken into categories and types. The various alert categories and types are:

Category	*Type*
Demand	Request Date Violation
Inventory	Item Maximum Level Violation
	Item Minimum Violation
	Item Safety Violation
Resource	Minimum Batch Violation
	Resource Utilization
Work Order	Operation Past Due
	Operation Fixed Date Violation
	Request Date Violation

When an alert is selected in the Model Workspace, the Alerts grid opens and the applicable tab is displayed.

A given alert category, such as Work Orders, may have more than one alert type. If the user selects a specific alert within the tree, the alerts within the tab are filtered to only that alert type. For example, if there are both Operation Fixed Date Violations and Operation Past Due alerts, and the user selects only the Operation Past Due alert in the model workspace, then only the Work Order tab is displayed with only those alerts that have Operation Past Due alert types.

The Resource category provides a Resource Utilization alert to indicate when resources which are not part of a bottleneck resource group have capacity violations (when capacity constraints are relaxed).

The Alerts Grid

The alerts grid displays information about any solver-based constraint violations. These violations do not necessarily mean that the schedule is infeasible but rather that a soft constraint can not be met. For example, the solver might be required to temporarily violate a minimum inventory level to meet demand. The alerts grid displays these violations in detailed tables so that you can investigate and improve the schedule.

The alerts grid has four tabs that correspond to the categories in the alerts tree: Demand, Inventory, Work Order, and Resource. The alert count is displayed next to the category description.

If desired, the alerts grid can be modified to view its contents better. By adjusting the alerts grid window size and column width, you can expand or reduce the amount of information visible.

Accessing Alert Messages

Use this page to access alert messages:
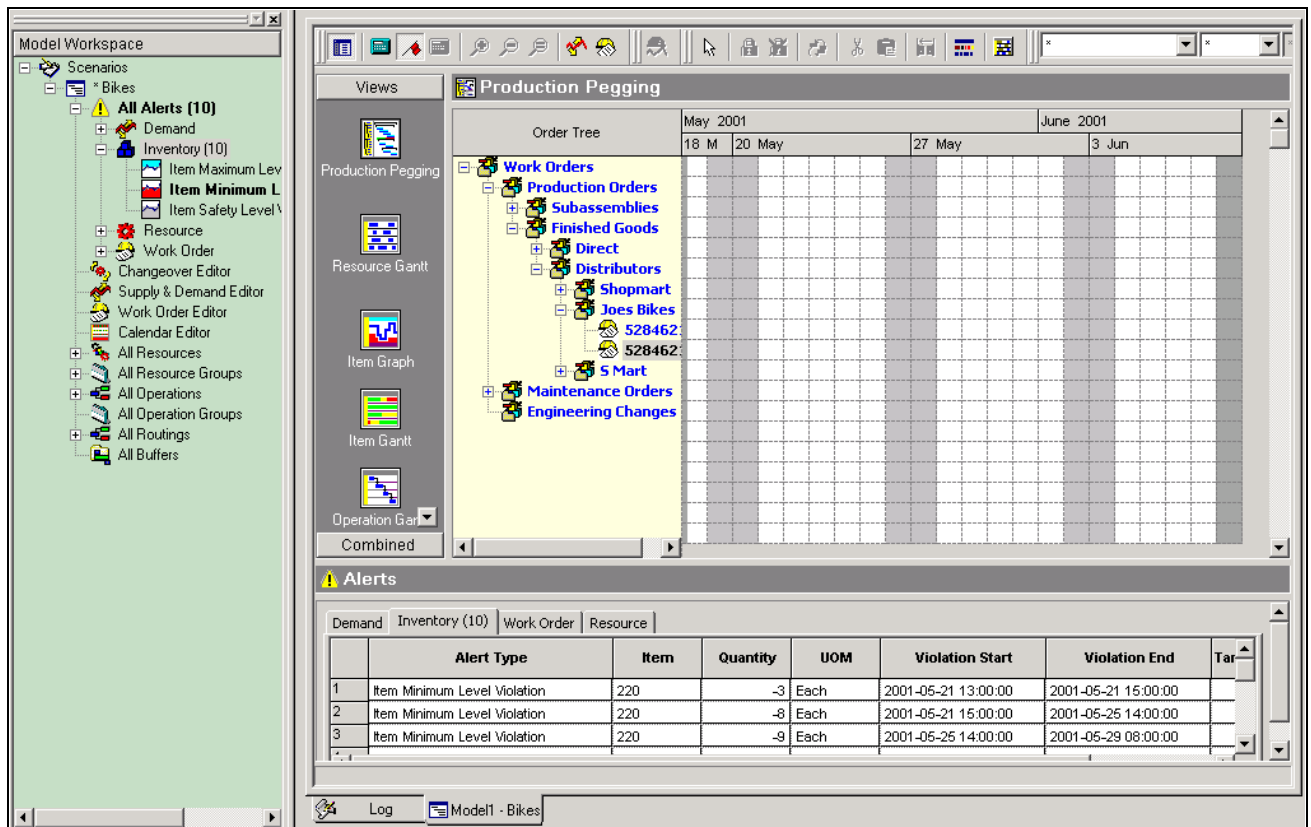
Accessing the Alerts grid

To access alert messages:

1. Access the Model Workspace.

2. Select one of these options:

    •   Expand the All Alerts tree option, which expands the tree to display the alert categories.

    •   Click the All Alerts text, which displays the Alerts grid below the view.

Navigating to Production Pegging Alerts

To navigate to the Production Pegging or Work Order Editor:

1. From Production Pegging Demand view, access the Alerts grid.

2. Right-click the alert row.

3. Select Drill down to Production Pegging or Drill down to Work Order Editor.

Filtering Alert Messages

To filter alert messages:

1. Access the Alerts grid.

2. Select the column heading of the field that you want to filter.

3. Select a field variable from the available options in the menu.

Working with KPIs

This section provides overview of KPIs and schedule comparison and discusses how to:

• Access the KPI view.

• Compare production schedules.

• Sort schedules by KPI.

• Reorder KPI groups.

• Hide KPI groups.

• Hide KPI columns.

• Reorder KPIs in groups.

Understanding Key Performance Indicators

After you create the initial schedule, it can be duplicated as many times as necessary, and you can make different changes to each version using manual scheduling. To determine which version of the schedule is the most cost-efficient or best fulfills customer orders, you can evaluate all the schedules against each other using KPIs. KPIs are objective measures that can provide you with important information about how the schedules perform from a customer service, material, cost, and manufacturing perspective. KPIs are available for solved schedules and are dynamically refreshed each time a schedule is solved or repaired.

KPIs are calculated values that enable you to evaluate schedules and select a schedule that best suits your specific criteria. The KPI view can be demand centric or work order centric. The KPIs are categorized into the following groups:

• Customer service

• Costs

• Manufacturing

• Materials

This page illustrates the KPI view:

KPI view

You can sort the schedules according to your choice of KPI, rearrange the order of the columns, or hide specific columns or groups from displaying in this view. The schedules that are displayed in the Key Performance Indicators windows can be returned to the original sorting order if you no longer want to sort by a specific KPI.

Column width and row height can also be changed. If you do not want to view a specific KPI or if you want to keep it from general view, you can hide that column. Hiding all of the columns in a group hides the group. If you prefer to evaluate schedules based on a specific KPI group such as Costs, you can reorder the KPI groups so that the group appears first in the grid.

Note. A number of the KPIs depend on pegging information for their calculations. Verify that you have specified the primary operations for each routing and the primary outputs.

Customer Service

The Customer Service group outlines metric information such as demand fill rates, number of stockouts, and so on.

This table explains how the columns are calculated and how to interpret them:

Column	*Description*
Order Fill %	This statistic is taken directly from the summary statistics in the Supply and Demand Editor, where these are available by folder, and by subfolder, which makes them even more useful. Order fill rate is calculated by dividing the total number of orders that have all their line items met on time by the total number of orders in the equation.
Line Fill	This statistic is taken directly from the summary statistics in the Supply and Demand Editor, where these are available by folder, and by subfolder, which makes them even more useful. Line fill rate is calculated by dividing the total number of on time line items by the total number of line items in the equation.
Unit Fill %	This statistic is taken directly from the summary statistics in the Supply and Demand Editor, where these are available by folder, and by subfolder, which makes them even more useful. Unit fill rate is calculated by dividing the total number of on time units ordered by the total number of units in the equation. For example, if you have two sales orders that have a total number of units ordered across various items of 10,000 and only 6,000 items are on time, then your unit fill rate is 60 percent.

Column	Description
Number of Stock Outs	This is the total number of times you drop below the greater of item minimum or zero in the horizon for all of your items.
Stock Out Time	This is the total amount of time in the horizon that the schedule is stocked out across all items. The time is expressed in the time unit in the scenario properties. See "Setting up the Scenario Properties - General Page," Configuring Scenario Properties, *Production Scheduling*

Costs

The Costs group provides various costs associated with the schedule such as changeover cost, manufacturing costs, and so on.

This table explains the Cost metric information presented in the KPI view:

Column	Description
Change Over Cost	This is the sum of costs for all scheduled changeovers in the schedule. Each changeover has a cost associated with it in the changeover matrix. Every time a changeover occurs, it is added to the total cost.
Inventory Cost	The inventory cost provides the total carrying cost offset by total stockout costs. You can specify carrying costs, stockout costs, and safety violation costs on an item. You can also specify your time unit. The time is expressed in the time unit in scenario properties. See "Setting up the Scenario Properties - General Page," Configuring Scenario Properties, *Production Scheduling* It is calculated as follows: sum of individual inventory costs, C_i, where C_i = (inventory carrying cost * area above zero) + (inventory stockout cost * area below zero) (areas from the Item Graph).
Manufacturing Cost	This cost reports the sum for all operating costs occurred in the schedule. It is calculated as follows: setup cost + (duration of each operation * (resource operating cost + operation operating cost).

Manufacturing

The Manufacturing group provides metric information that includes plant throughput and resource utilization.

This table explains the Manufacturing metrics in the KPI view:

Column	Description
Change Over Time	This is the sum of time for all scheduled changeovers. The time is expressed in the time unit in scenario properties. See "Setting up the Scenario Properties - General Page," Configuring Scenario Properties, *Production Scheduling*
Machine Utilization %	This is the average for all reusable resources of run time divided by up time. Run time is the sum of all operations that are run on this machine. Up time is the schedule horizon minus the sum of delay and down calendars on the resource.
Throughput (Can be viewed by shift, day, week, or month.)	This is the sum of the quantity of all demands divided by plant cycle time. Plant cycle time = End Time of last scheduled operation − Start Time of earliest scheduled operation (of all operations in the schedule). For example, if it takes 10 days from the start of the first operation to the end of the last operation to produce 10 units, the shift throughput is 3.33 units, daily throughput is 10 units, weekly throughput is 70 units, and monthly throughput is 280 units.
Manufacturing Cycle Time	The average amount of time that it takes to manufacture items during the schedule horizon. For each demand routing instance = (End Time of Latest Operation - Start Time of First Operation) / Number of Units of Item Produced. The time is expressed in the time unit in scenario properties. See "Setting up the Scenario Properties - General Page," Configuring Scenario Properties, *Production Scheduling* For example, if time is set to days and you can produce 1 unit in 24 hours for an operation, you have demand for 1 unit. If you produce 1 unit, this value is displayed as *1*. If you change the time unit to hours, it is displayed as *24*.

Materials

The Materials group provides various metric information that outlines inventory turns, substitutions, and so on.

This table explains the Materials metrics in the KPI view:

Column	Description
Inventory Turns	Inventory turns are an average for all items that have demand. Inventory turns are calculated as follows: total demand quantity / average item level in the horizon.

Column	Description
Periods of Cover Which can be viewed by shift, day, week, or month.	Average of Pi, where Pi is the number of periods before a demand that is not covered by the current inventory level at the start of a period.
Number of Substitutions	The total number of times a lower priority item is selected from an item set

Understanding Schedule Comparison

Comparing schedules is ideal if you want to make a change to an existing schedule and assess its impact to the original schedule. You begin by renaming the original schedule Baseline Schedule. A duplicate of this schedule can be made as many times as necessary. Then, as changes are made to each of the duplicate schedules and repaired, rename the schedules with more descriptive names. With all the schedules contained in a schedule folder, you can then compare these schedules.

Compare the KPIs for a schedule against a baseline schedule to determine the net difference in customer service between schedules. You can also select the colors used to highlight the baseline and comparison schedules in the Key Performance Indicator window.

This page illustrates that the baseline schedule is in blue and the schedule titled, "Run Overtime" is in yellow:

KPI comparing schedules

By comparing and refining schedules in the KPI window, you can continually improve the production planning process and ensure that your enterprise has the materials available to meet the requirements in the production plan. The KPIs displayed in the Key Performance Indicators window are updated automatically when you solve your schedule.

The Demand Details pane highlights changes in demand fill data between a schedule and a baseline schedule. The following table describes the demand fill data that is included in the Demand Details pane.

Field	Description
Demand	The demand code for the schedule and the associated line item.
Item	The item that is scheduled for production.
Line Item	The line item number of the demand.
Order Class	The order class of the line item.
Priority	The demand priority of the item that is scheduled for production.
Customer	The customer that is requesting the item.
Change in Makespan	The change in the duration of the demand, including associated line items and routings. The value is expressed in days and hours. A negative value indicates a shorter makespan.
Change in Lateness	The change in lateness between the baseline schedule and the schedule used in a comparison. The value is expressed in days and hours. A negative value indicates that the lateness of the demand is reduced.
Request Date	The request date indicated by the baseline schedule. For new demand, the request date is taken from the comparison schedule.
Available Date	The available date indicated by the baseline schedule.
Available Date (Comparison Schedule)	The available date indicated by the comparison schedule.

Accessing the KPI View

Use this navigation path to access the KPI view:
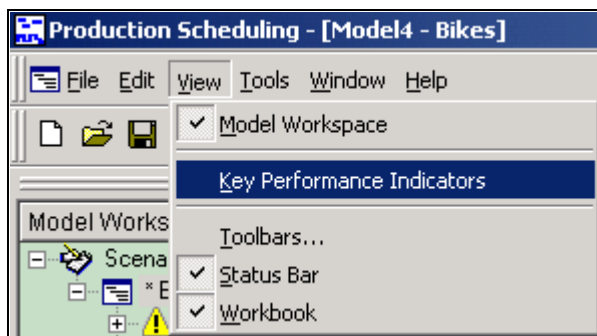
Accessing the KPI view

To access the KPI view:

1. From the menu bar, select View.

2. Select Key Performance Indicators.

   The KPI view opens.

3. View the schedules in the following KPI categories:

   • Customer Service

   • Materials

   • Costs

   • Manufacturing

Comparing Production Schedules

To compare schedules:

1. Access the Key Performance Indicators window.

2. In the Schedule Comparison pane, select a schedule to use as the baseline from the Baseline Schedule menu.

3. Select a schedule to compare with the baseline from the Compare to menu.

   The baseline schedule and the schedule used to compare with the baseline schedule are highlighted in the Key Performance Indicators window. You can change the highlighting colors by selecting Customize Colors on the toolbar.

4. From the Demands menu, select one of the following values to display changes to demand in the new schedule:

   • Was On-time - Now Late

   • Was Late - Now On-time

   • Was Late - Less Late

   • Was Late - More Late

      An asterisk (*) displays all demands that exist in both schedules.

5. Select Include New Demands In Comparison Schedule to include any new demands that were not included in the original schedule.

   Metrics for the baseline schedule and the comparison schedule are displayed in the Demand Details pane.

Sorting Schedules by KPI

To sort schedules by KPI:

1. Access the Key Performance Indicators window.

2. Click the KPI column heading for the KPI in which you want to sort.

    This action sorts your schedules in ascending order.

3. Click the KPI column heading again to sort your schedules in descending order.

4. Click the top left blank column heading to return to the original order.

Reordering KPI Groups

To reorder KPI groups:

1. Access the Key Performance Indicators window.

2. Right-click in the top left blank column heading and select Hide Groups.

3. In the Shown Groups field, select the group you want to move to a different position.

4. Click either the Move Up or Move Down button to change the position of the group relative to the other groups.

5. Click OK.

Hiding KPI Groups

To hide KPI groups:

1. Access the Key Performance Indicators window.

2. Right-click in the top left blank column heading and select Hide Groups.

3. On the Show/Hide Grid Columns window, in the Shown Groups field, select the group that you want to hide.

4. Click the button with the double arrows pointing to the right to move the group to the Hidden Groups field.

5. Click OK.

Hiding KPI Columns

To hide KPI columns:

1. Access the Key Performance Indicators window.

2. Right-click the KPI column that you want to hide and select Hide Columns.

3. In the Show/Hide Grid Columns window, in the Shown Columns field, click the KPI that you want to hide.

4. Click the button with the double arrows pointing to the right to move the KPI to the Hidden Columns field.

5. Click OK.

Reordering KPIs in Groups

To reorder KPIs in groups:

1. Access the Key Performance Indicators window.

2. Right-click in the column heading and select Hide Columns.

3. In the Shown Columns field, select the KPI that you want to move to a different position.

4. Click either the Move Up or Move Down button to change its position relative to the other KPIs in the group.

5. Click OK.

Chapter 10

Manual Scheduling

This chapter provides an overview about solved schedules and discusses how to invoke manual schedule changes.

Understanding Solved Schedules

This section provides an overview of production schedule adjustments and discusses:

- Manual scheduling.

- Operation adjustments.

- Resource adjustments.

- Resource idle time removal.

- Batchable resources.

- Resource calendar event modification.

- Schedule repair.

Understanding Production Schedule Adjustments

After an initial schedule has been created, some schedule adjustments may be required to make improvements to the schedule. Production Scheduling provides the necessary tools to analyze and modify the schedule.

Any item or demand violations that are identified during the solve are displayed in the alerts grid. The suspected operation that is causing the violation is listed in the alerts grid and can be expanded and highlighted within the Resource, Demand, and Operation Gantt schedule views. Once the operation is identified, then the root cause of the violation can be investigated and corrected.

A wide range of changes can be made to the schedule after its initial solve, including changes to the following:

- Operation durations.

- Operation start and end times.

- Assignment of operations to resources.

- Sequence of operations.

- Downtime.

- Demands.

- Quantities.

Operations can also be fixed to run at a specific time. Changes that are made to the schedule are highlighted for easy identification.

Manual Scheduling

Manual scheduling enables you to make manual changes to the schedule so that you can respond to new information from the shop floor, alerts generated by the solver, or changes to the maintenance calendar.

This list shows the wide range of changes that can be made to a schedule after its initial solve:

- Add new demands.

- Modify operation durations.

- Modify start and end times.

- Change assignment of operations to resources.

- Resequence operations.

- Modify delay and downtime events in the Resource Calendar.

- Remove idle time between operations.

- Adjust resource capacities.

- Adjust inventory levels.

- Adjust demand quantities.

- Adjust dates.

Operations can also be fixed to a run at a specific time. Demand can be modified locally for the active schedule.

The toolbar enables you to make some of the manual changes to the schedule and run a repair solve to update the schedule.

This example illustrates manual scheduling changes in the Resource Gantt:
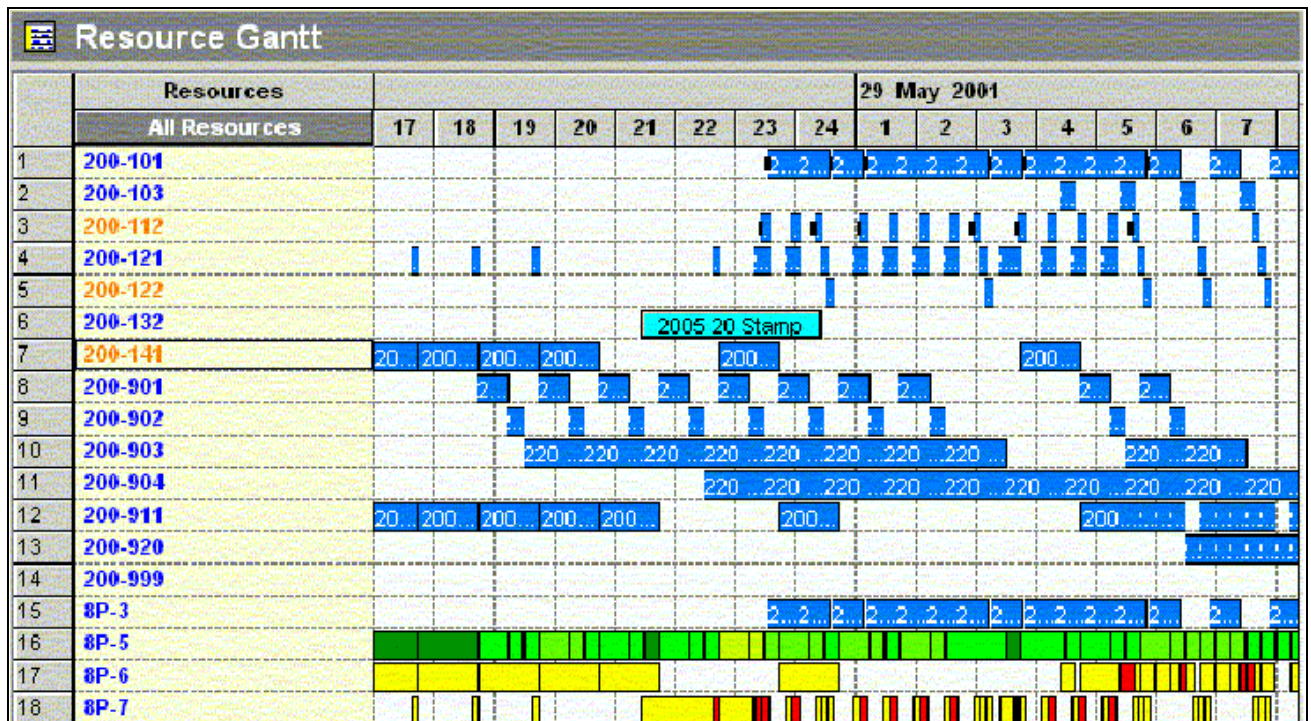
Manual scheduling changes

Manual scheduling changes are highlighted in light blue for easy identification, as displayed with the operation 2005 20 Stamp in the example. In comparison with the original solve, the repair solve focuses only on those areas that are changed in the schedule when producing a new schedule. In many cases, this action can represent a significant time savings. The manual scheduling changes are logged to the log window for later reference. After the repair solve is complete, manual scheduling changes are highlighted in orange in the Resource Gantt.

Operation Adjustments

Operation adjustments can be made by:

• Using cut and paste commands.

• Using the Enforce Horizon Start option.

Using Cut and Paste Commands

Production Scheduling uses the following rules when pasting operations:

• In the first scenario, an operation is pasted into idle time.

In this case, the operation is pasted at that point.

• In the second scenario, operation A is pasted onto operation B.

In this case, operation A is pasted into idle time immediately following operation B. If another operation immediately follows operation B, then operation A is pasted between these two operations.

When you run a repair, the solver tries to move the operation to the exact point specified. If constraints prevent the operation from moving to the specified start time, then the solver will try to reschedule the operation later during the idle time. If pasting the operation results in an infeasible schedule, then the system cancels the paste command and moves the operation back to its original location.

If you select the Fix option in the Paste Operation dialog box, then the operation must start at the time specified in the Start Time field. Pasting a fixed operation into idle time or on top of an existing operation might result in an infeasible schedule.

Note. When you cut and paste multiple operations with at least one operation that has the same target resource as the original resource, the system disables the Ignore paste start time and sequence option.

When an operation is pasted into its new location, the solver can resequence some of the adjacent operations to accommodate this change (for example, if the idle time that the operation being pasted into is not long enough). Production Scheduling gives you the option of performing either multi- or single-stage resequencing.

Single-stage resequencing honors any solver-imposed constraints on the operation. If the resulting paste violates these logical constraints (for example, if an operation that produces an item is moved after the operation that requires the item), then the system does not move the operation.

Multistage resequencing enables the solver to move all of the operations that constrain the pasted operation. For example, assume that an operation is pasted before the operation that produces an item that it requires. In this instance, multistage resequencing enables the solver to move the affected downstream and upstream operations.

The duration of an operation can be quickly modified to display its impact on an existing schedule. If any resources are affiliated with the changed operation, they are also modified to stay consistent with the new duration.

By changing an operation's start and end times you can manually reschedule an operation within the schedule. An operation's start and end times can be changed from the Demand, Operation, or Resource Gantt views.

The Start Time and End Time fields represent the actual times when the operation instance is scheduled. The Earliest Start Time and Latest End Time fields represent the time frame in which this operation instance could be scheduled without definitely causing infeasibilities. You should select a start time and end time within these parameters to render a feasible schedule.

The start and end times of merged operations can be modified as a single entity. After changing a time, the groups of merged operations are each shifted accordingly. Note that the Run Time, Delay, Earliest Start, Latest End, and Elapsed Time properties cannot be changed for merged operations.

When the start or end times are changed, the duration remains unchanged and the operation is moved to the chosen date and time. To change an operation's duration, modifying the value in the Duration field.

Note. You can also modify start and end times by cutting and pasting operations from the Operation and Resource Gantt views.

Using the Enforce Horizon Start Option

When schedule backlog occurs, you can elect to use the Enforce Horizon Start option from the Resource Gantt view. The user selects the applicable resource, performs a right-click, and selects the Enforce Horizon Start option. The user then performs a repair solve or a solve.

When the enforce horizon start option is used, all operations for that resource are pushed forward to the start of the horizon. Operation sequence is maintained and precedence constraints for upstream and downstream operations are respected.

When selecting the Enforce Horizon Start option on the Resource Gantt view:

• The resource attribute persists and is added to the data model.

    For example, if a user selects this option, it is persisted on the resource properties and displays a check mark beside the text in the right-click menu. This notifies the user that the flag has already been selected if the user right-clicks on a resource again.

• Enables the Enforce Horizon Start option on the Resource Properties - Availability tab.

• The fixed timefence changes to a cross hatched pattern for machines, crews, and tools.

Users can highlight many resources at once and enable this flag for all highlighted resources. However, if a resource that is highlighted already has this option selected, its state is reversed.

After selecting the Enforce Horizon Start option, if the user clicks the Repair solve button, all operations are moved forward in time, keeping the same sequence but beginning at horizon start. When this sequence is performed:

• Operations running in the past start at the beginning of the horizon.

• Any adjacent operations are pushed forward to accommodate the operations moved forward in time.

• The sequence of operations on the resource are not changed.

• Any upstream or downstream operations of an operation may need to be pushed forward if precedence constraints dictate.

After selecting the Enforce Horizon Start option, if the user clicks the cold solve button, the solver resolves resource conflicts and enforces horizon start in a one-step process. Operations that appear to be running behind are moved to the horizon start. The difference from a repair solve is that the user is not alerted that some operations are running in the past. By selecting the Enforce Horizon Start option and running a solve, this step is performed in one task.

In some cases, the operations running behind are one of several operations that are contained within a routing. Because routings have precedence constraints between operations, moving operations from the past to the horizon start may have consequences on other operations in the routing. For example, an operation is the first operation of a routing and it is pegged to a downstream operation. If Enforce Horizon Start is selected and a repair solve is performed, the pegged operation is moved forward in time to avoid a precedence infeasibility. Similarly, if an operation has been scheduled in the past due to a date conflict with another resource and upstream supporting operations have not yet run, those operations should be pulled forward to start just in time (JIT) to the placement of the operation moved to the start of the horizon.

Resource Assignments

For operations that use sets, Production Scheduling tries to solve the schedule based on the preferences that you apply to each set, while considering all constraints and priorities in the model. For instance, if you have specified one crew set with the option of using either Crew A or Crew B, the solver creates a schedule trying to use your preferred crew, as indicated in the properties of the operation model. However, after creating the schedule, you might want to change the resource actually used by the operation.

Note. You cannot manually change resource assignments for merged operations.

Resource Idle Time Removal

Although the JIT build method is often a good strategy to avoid excessive prebuilding, it can sometimes cause a fragmented schedule. Running machines nonstop for a period of time is often more efficient, especially if you also remove idle time from the schedule. You can use the Remove Idle Time command from the Resource, Demand, and Operation Gantt views to:

• Remove idle time between two or more selected operations.

• Remove a resource's idle time between an operation and a specified time within the schedule horizon.

• Remove all idle time on a selected resource.

Note. You can use the Remove Idle Time command on a multicapacity resource only from within the Operation Gantt view. If a multicapacity resource is selected on the Resource or Production Pegging views, the Remove Idle Time button is disabled.

Resource constraints can prevent the solver from removing all idle time. For example, assume that items that are required by an operation must be in inventory or a crew must be available to run a machine at the new time slot. In cases where operations cannot be moved, the solver removes as much idle time as possible.

Batchable Resources

Users can manually move operations between batches or manually move an entire batch. For example, if a sequence on a resource is scheduled to be 10 hours at 500 degrees followed by 12 hours at 400 degrees, the user can switch the order. This can be accomplished by using the cut and paste, drag and drop, or resequencing window in the Resource Gantt view.

When manually scheduling, you can move specific operations between batchable resources, known as off-loading, to fill up capacity. If you have attributes scheduled on a batchable resource that are different, but they are close enough to be able to run the operation with different attributes, and the final product quality is not affected. For example, some aluminum blocks can be processed at a particular temperature for a given time range. For example, 19 hours at 600 degrees. However, the quality of the final alloy is virtually the same if the job is processed for 20 hours at 590 degrees. Production Scheduling in itself is not able to understand alternate attributes, but you can manually schedule operations between resources, or the same resource, running a different attribute family type.

See Also

"Batchable Resources," Creating Production Scheduling Models, Defining Resources, Understanding Resources, *Production Scheduling*

Resource Calendar Event Modification

You can add, delete, or modify calendar events in the Resource Calendar. Alternatively, you can delete or modify calendar events in the Resource Gantt view. This function enables you to analyze the effect of adding another shift or resource to the schedule. After repairing the schedule, Production Scheduling updates the schedule to reflect the new events.

Calendar events can be added directly from the Production Pegging and Resource Gantt views. This feature is helpful when fine-tuning the schedule once it has been solved. These changes are only reflected in this schedule and are not applied globally to the Calendar editor.

Calendar events can be deleted directly from the Resource Gantt view, which is helpful when fine-tuning the schedule, once it has been solved. These changes are only reflected in this schedule and are not applied globally to the Calendar editor.

You can delete calendar events for a resource as required. After repairing the schedule, Production Scheduling updates the schedule to reflect the deleted events. When you delete a calendar event from a resource, the system does not remove the event from the parent calendar.

Note. When an event has been added in a schedule, the resource reassignment is not applied until the schedule is repaired. When a delay event has been deleted in a schedule, the resource reassignment is not reflected until the schedule is repair solved.

Schedule Repair

Repairing a schedule enables you to implement changes that are made to the schedule. To minimize the effect on the schedule, the repair solve modifies only the schedule for the resources and time periods that are specified. Depending on the changes that are made, a repair usually takes less time to complete than the initial solve.

If a change is made to the schedule that cannot be incorporated with a repair solve, then the Repair button is disabled. If a manual scheduling change results in an infeasible schedule, all of the changes since the last successful solve are reversed. After the repair solve is complete, manual scheduling changes are highlighted in orange in the Resource Gantt view.

Performing an Undo for Manual Scheduling

Manual Scheduling in Production Scheduling includes a variety of activities such as:

• Cut/Paste.

• Resequence.

• Remove idle.

• Change durations of operation instances.

• Change start or end times of operations.

• Change calendar events in the Gantt charts (add/remove/change).

- Change resource assignment (for example, manually offload).

- Fixing/Unfixing operations (single or multiple operations).

- Change lead times for Purchase Order Recommendations.

Following each of these actions, there is a repair-based solve action to institute the change and "repair" any upstream/downstream ramifications this change may have caused. However it is possible that the user may not accept the repaired schedule or simply made a mistake when implementing this change. This enhancement to Production Scheduling provides the ability to "Undo" these changes. The Undo action reverts back to the previous "solve state" prior to the change. The table below provides examples using the Undo feature.

Action	*Result*
{manual change}	UNDO	Will undo the manual change.
{manual change}	Repair	UNDO	Will undo the repair and manual change.
{manual change}	{manual change}	UNDO	…	Will undo both manual changes.
{manual change}	{manual change}	Repair	UNDO	Will undo the repair and manual changes.
{manual change}	Repair	{manual change}	Repair	UNDO	Will undo the 2nd repair and 2nd manual change.
{manual change}	Repair	{manual change}	Repair	UNDO	UNDO	Will undo the both repairs and both manual changes.

Undo also applies to another new feature in Production Scheduling (Resource Business Objectives). Included in this feature is the ability to solve groups of resources in a "Stage" and perform manual scheduling activities between stages (for example, At a breakpoint) in the solve process. Undo can be performed within this context as well. This includes the ability to undo manual changes in the stage itself as well as to completely undo a given "stage" of the solve process - upon which the solution will revert back to the previous solved "stage". The table below provides examples using the Undo feature.

Action	*Result*
{solve stage	breakpoint}	manual change	Repair	UNDO	Will undo the manual change.
{solve stage}	{solve stage	breakpoint} UNDO	Will undo the 2nd solve stage – result will revert back to the first solve stage.

Invoking Manual Schedule Changes

This section lists prerequisites and discusses how to:

- Move operations within a routing.

- Modify operation durations.

- Modify operation start and end times.

- Modify operation resource assignments.

- Fix operations.

- Resequence operations.

- Remove resource idle time.

- Repair a schedule.

- View manual schedule changes in the Gantt view.

- Add calendar events in the resource calendar.

- Enforce horizon start time.

Prerequisites

Before you begin to make manual scheduling changes, ensure that you meet the following prerequisites:

- Solve your schedule.

- Use either the Demand, Resource or Operation Gantt view.

Moving Operations Within a Routing

Access the Operation Gantt, Resource Gantt, or Production Pegging view.

To move operations within a routing:

1. Select the operation that you want to move.

   To select multiple operations, hold down the left mouse button and surround the operations. The system displays the selected operations in red.

2. On the Schedule view toolbar, select Cut Operation from the menu bar or right-click your mouse.

   The cursor changes to a horizontal scroll bar that displays the current date and time. Allowable destination resources for the cut operation appear highlighted in blue.

   Note. If you select operations across multiple resources that are not able to run on the same machine, the Cut Operation button is disabled.

   ---

3. Click the location corresponding to the date and machine where you want to move the operation.

   The selected time appears in the status bar.

4. Right-click and select Paste.

   The Paste Operation dialog box appears and displays the selected start and end times.

   You can set the operation to start at the earliest available start time by selecting the Earliest Start Time check box or you can manually adjust where you want to paste the operation by changing the Start Time and End Time fields.

5. Select one of the following options to indicate how you want to resequence adjacent operations:

Multi-stage Resequencing	Select this option to enable the solver to move all of the operations that constrain the pasted operation.
Single-stage Resequencing	Select this option to honor any solver-imposed constraints on the pasted operation.

6. To off-load an operation from one machine to another, select Ignore paste start time and sequence.

   When you run a repair, the solver schedules the moved operation. This option is only enabled if the machine that you are using is a part of a resource set.

   Note. When you enable the Ignore paste start time and sequence option, the Start Time, End Time, Duration, and Fix fields are disabled.

7. Select OK.

Modifying Operation Durations

Access the Operation Properties window. Select the Times tab.

To modify operation durations:

1. In the Run field, enter the duration of time that you want to use.

   The End Date field is adjusted to reflect the new operation duration.

2. Click OK.

   Note. When changing operation duration, the new time must fall within the bounds set by the Earliest Start Time and Latest End Time fields.

Modifying Operation Start and End Times

Access the Operation Properties window. Select the Times tab.

To modify operation start and end times:

1. In the Start Time field, enter the time when you want the selected operation to begin.

2. In the End Time field, enter the time when you want the selected operation to end.

3. Click OK.

The change in start and end times is displayed on the schedule. Any scheduled operations to be executed after the end of the changed operation are modified as well.

Modifying Operation Resource Assignments

Access the Operation Properties window. Select the Resources tab.

To modify operation resource assignments:

1. Select the code field for the set with resources that you want to change.

2. Select the down arrow to view the resource options.

3. Select the resource that you want from the list.

4. Click OK.

Fixing Operations

Access the Schedule view.

To fix operations:

1. Select Multiple Selection from the Schedule toolbar.

2. Click once on each operation that you want to fix.

    The operations change from blue to red.

    If you select an operation by accident and want to deselect it, click the operation once again. It reverts to its original blue color.

3. Select Fix Operations from the Schedule toolbar.

    The fixed operations change from red to green.

4. To undo the fixed operations, select Unfix Operations from the Schedule toolbar.

Resequencing Operations

Access the Production Pegging or Resource Gantt view.

To resequence operations:

1. Select the operations that you would like to resequence.

2. Select Resequence Operations.

    The Operation Resequencing dialog box appears, listing all of the operations that are selected for resequencing in chronological order.

3. From the left row headings, select one of the operations.

   The Move Up and Move Down arrow options at the top of the dialog box activate.

4. Click the appropriate button to move this operation up or down.

5. Repeat steps 3 and 4 for each operation that you want to move.

6. Click OK when complete.

Removing Resource Idle Time

Access the Production Pegging, Operation Gantt, or Resource Gantt view.

To remove resource idle time:

1. Select one of the following options:

   • On either the Resource or Production Pegging view, use the multiselection functionality to select two or more operations on a single-capacity resource.

   • On the Operation Gantt view, use the multiselection functionality to select an operation type.

2. Select Remove Idle Time.

   All of the operations within the selected range are moved to an earlier time so that all of the idle time between them is removed. If more than two operations are selected, the solver removes idle time between the earliest and the latest selected operation.

   Note. If the operations that you select do not run on the same resource, the system disables the Remove Idle Time command.

3. Select Repair to repair the schedule.

   Note. After a repair, the solver attempts to move all of the selected operations to remove excess idle time. However, resource constraints (for example, an operation running within a fixed timefence, or an operation that is waiting for inventory) can prevent the solver from removing all idle time. In cases where operations cannot be moved, the solver removes as much idle time as possible.

Repairing a Schedule

Access the Schedule window.

To repair a schedule:

1. Complete manual scheduling adjustments.

2. Select Repair.

The schedule is refreshed with the incorporated changes. If the repair results in an infeasible schedule, a message appears on the status bar indicating that the schedule is infeasible. The schedule returns to the previous solved state.

Note. Some manual scheduling changes require that you perform a solve. For example, schedules containing modifications to routing and operation properties, work orders, and supply and demand data cannot be repaired and require a solve. In these cases, the Repair tool becomes unavailable for input.

Viewing Manual Schedule Changes in the Gantt Views

To view manual scheduling changes:

1. Access the Production Pegging, Operation Gantt, or Resource Gantt view.

2. Select Highlight operations changed by Repair on the Gantt toolbar.

The operation that you manually changed before a repair is highlighted in red. Related operations in the routing that have been changed by the solver during a repair are highlighted in orange.

Adding Calendar Events in the Resource Calendar

To add calendar events:

1. Access the Resource Calendar.

2. Click the down arrow under Resources, and select the resources or groups that you want to view.

3. Click a resource, and then click the Add Event button.

4. Click one of the following options to specify the event type:

Up Time	All machines and activities are available.
Down Time	Actual downtime when all machines and activities are stopped. This option imposes significant constraints and should be used infrequently. Use the Delay Time option whenever possible.
Delay Time	A delay when machines and activities are temporarily suspended.

5. Complete the following fields:

Start Time	The date and time when the event begins. Alternately, you can leave the original date. When the end time and the duration are entered, the start time is automatically calculated.
End Time	The date and time when the event ends. Alternately, you can leave the original date. When the start time and the duration are entered, the end time is automatically calculated.
Duration	The event length in seconds, minutes, hours, or days. If the start time is specified, the end time is automatically filled in, based on the duration. Similarly, if the end time is specified, the start time is filled in automatically.
Recurring Event	Select this option if you want the calendar event to repeat. Select one of the following options in the Recurrence Pattern section: Daily, Weekly, Monthly, or Yearly.

6. Click one of the following options in the Recurrence Limit section:

No end date	Select this option if the event occurs throughout the calendar horizon.
End after occurrences	Select this option and select the number of occurrences if the event occurs for a specified number of times, and then stops.
End by	Select this option and select an end date if the event ends by a specified date.

7. Click Add.

   Production Scheduling adds the calendar event to the Resource Calendar.

Enforcing Horizon Start Time

There are two methods to enforce horizon start time:

The first method to enforce horizon start time is:

1. Access the Resource Gantt view.

2. Right-click on a resource.

3. Select the Enforce Horizon Start option.

The second method to enforce horizon start time is:

1. Access the Resource Gantt view.

2. Right-click on a resource.

3. Select the Properties option.

4. Select the Availability tab from the resource properties window.

5. Select the Enforce Horizon Start option.

Chapter 11

Publishing Data

This chapter provides an overview about publishing Production Scheduling data, Production Scheduling published reports, and the web dispatch list, and discusses how to:

- Publish schedules.

- Publish a schedule to Oracle E-Business Suite

- Export schedule data.

Understanding Production Scheduling Published Reports

You can publish data from Production Scheduling in any of the following formats:

- CSV

- HTML

- XML

CSV File Reports

CSV files can be read into an ETL tool of your choice for transformation in preparation for import into other systems. When you publish in .csv file format, you can generate one or more of the following reports:

- Operations Report

- Resource Schedule

- Demand Fulfillment

- Key Performance Indicators (KPIs)

HTML Dispatch Report

HTML file format generates an HTML Dispatch Report that can be published directly to the shop floor.

Oracle E-Business Suite Integration

Oracle E-Business Suite Integration file format generates XML files, which are used to release orders to E-Business Suite.

XML File Reports

The standard XML files, which are exported out of Production Scheduling are:

• ProductionSchedule.xml

• PurchaseOrderRecommendations.xml

• WorkOrderSchedule.xml

XML files can be compressed when publishing, which reduces the size of the files created.

Understanding the Web Dispatch List

This section discusses:

• Publishing options.

• Index page.

• Resource details page.

Publishing Options

Production Scheduling enables you to send production schedule data to the production line using the web dispatch list. The web dispatch list eliminates the need to create custom reports and distribute them manually. Any production schedule that has been generated can be dispatched to a specific directory and posted online. Production personnel can access the web dispatch list to determine which scheduled operations are in the queue for a specific date range.

When you publish the web dispatch list, you specify a start date and a publishing horizon. The start date is the first day that you want to publish scheduled operations for resources in your supply chain. The publishing horizon is the length of the web dispatch list in days. The publishing horizon determines the size of the web dispatch list, according to your requirements. You can select publishing options to specify what data should be published in the web dispatch list. The following publishing options determine the level of data that is published on the resource details page:

Publishing Option	Data Published
Operation Description	A description of the operation that runs on the specific resource.
Start Time	The operation start time from the data model.
End Time	The operation end time from the data model.
Item	The item that is produced using this specific operation.
Item Description	A description of the item that is produced using this specific operation.

Publishing Option	*Data Published*
Quantity	The quantity of the item produced by this specific operation.
Work Order Number	Pegged work orders for this specific operation, indicating the work order that requires this operation.
Demand Number	Pegged demands for this specific operation, indicating customer demands that require this operation.
Demand Due Date	The due date for the customer demand.
Customer Name	The name of the customer who is requesting the item produced by this operation.
Customer Number	The customer code for the customer who is requesting the item produced by this operation.

Index Page

After the production schedule data has been published to the web dispatch list, the index page displays scheduled resources that are required for the current day. The dispatch list contains each resource that is required for the current schedule and a description of each resource. Resources are organized by the resource categories that are defined in the data model.

The index page is dynamically generated when you publish the web dispatch list in the desktop. You can define the location of the index page when you configure system options during initial system configuration. You must publish the dispatch list for each schedule that you generate or after you make manual scheduling changes and re-solve the schedule.

Resource Details Page

From the index page, scheduled resources can be selected in order to view detailed duration, item, quantity, pegging, and customer data for each operation that runs on that specific resource. Using this information, your production personnel can view an accurate schedule of production activities for a specific scheduling horizon, enabling them to meet the scheduled demands more efficiently.

Publishing Schedules

This section provides an overview of publish profiles, an example of a publish horizon, and describes how to:

•   Create a publishing profile.

•   Select reports.

•   Approve a schedule.

•   Publish a schedule using publish profiles.

- Edit an existing publishing profile.

- Delete an existing publishing profile.

Understanding Publish Profiles

Before you can publish a schedule, you need a publish profile and you must approve the schedule that you want to publish. Basic publishing profiles are provided with Production Scheduling; however, you may want to create personalized profiles that are more suited to your specific requirements.

When creating a publish profile, you determine general options such as the publish horizon, the format in which the data is to be published, and the output directory. Your data selection is based on the selected format. Data selection options are individual reports, if the output is to be compressed, or report headings.

As part of the Publish process, you have the ability to publish the Production Scheduling model, or selected objects of the data model, in XML format. For example, you can publish changeover rules, groups, crews, machines, tools, operations, and so on. This is useful when Production Scheduling is the system of record for any data that is not available in the ERP system. The next time that a model is imported into Production Scheduling, integration processes can be extended to reference the various data model elements that were exported from the previous scheduling session, which could amend to the new model prior to import.

When creating a publish profile, use the Model tab to define which data objects to publish. When the Publish Process is called, whether in batch or interactively, the selected data is exported to a user-named XML file. Interactively, the published model is based off of the approved schedule. Because the schedule publish process batch file points to a named scenario, you can export the model from any given scenario that may or may not be the approved schedule.

Example of a Publish Horizon

For example, assume the following three operations are part of a routing to make item *A*. A publish timefence of 2 days is in place. This diagram illustrates the scenario:
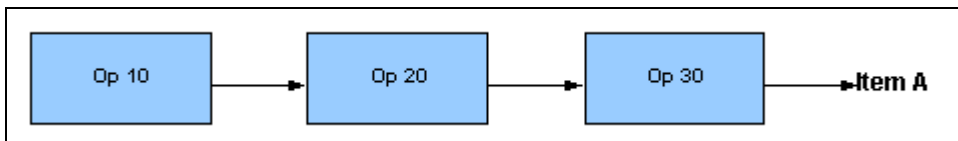
Publish timefence with pegging

The following logic indicates the data that is published when a timefence is being utilized:

- If any of Operation 10, 20 or 30 is inside the publish timefence, the routing and operations for those operations are included in the publish of the production schedule.

- If the routing instance crosses the timefence, the operations are included as part of the publish.

  If none of the operations associated with the routing cross or start within the timefence, then the operations are not published

- In terms of the pegging reports produced from Production Scheduling, only those demands which have operations that take place inside the timefence will have their pegging details published.

    If a series of operations do not have any activity within the timefence, the pegging details are not published.

Creating a Publishing Profile

To create a publishing profile:

1. Access the Production Scheduling desktop.

2. From the File menu, select Publish.

3. In the Publish window, click Add. The New Profile dialog appears.

4. Complete these fields on the General tab:

Profile Name	Specify a unique name for the profile. When you publish a schedule, you are asked to select a profile. The profile that you select will follow the criteria you set out in this procedure.
Publish Horizon	Specify the number of hours, shifts, days, or weeks for the profile. The publish timefence within a publish profile specifies a period of time in which data is published from Production Scheduling. See Example of a Publish Horizon.
	See "Example of a Publish Horizon," Publishing Data, Publishing Schedules, *Production Scheduling*
Format	Select *CSV, HTML,* or *XML.* This is the format in which your data is published.
Output directory	Specify the directory where the system saves the published data. You can explicitly specify the location of the directory or click Browse to browse for a directory location.
Script	Specify a script that is initiated after the data is published. You can explicitly specify the location of the script or click Browse to browse for a script. This field is optional.

5. Click Apply and then OK.

6. Select the Data Selection tab.

7. Depending on your selection in the Format field on the General tab, determine these additional options:

See "Selecting Reports," Publishing Data, Publishing Schedules, *Production Scheduling*

- If you select *XML,* select the reports that you want to publish, and select Compress Output if you want to compress the output report.

   The output is compressed into a GZIP file and .gz is appended to the file name. For example, if you compress the Production Schedule report, the file is ProductionSchedule.xml.gz.

- If you select *CSV,* select the reports that you want to publish.

- If you select *HTML,* select the boxes beside each heading that you want to appear in the Dispatch Report.

8. Click Apply and then OK.

9. Select the Model tab and complete these fields:

Export Model	Select to indicate that you want to export the selected data elements when you publish the file specified in the File Name field.
File Name	Select the user defined XML filename.
Model Data Selection (check boxes)	Select the data elements that you want to publish.

Selecting Reports

This section describes the reports that are available on the Data Selection tab when you select CSV, XML, or HTML as your publish format.

See "Creating a Publishing Profile," Publishing Data, Publishing Schedules, *Production Scheduling*

Operations Report

The Operations Report provides the actual start time, end time, and quantity of each item that is produced or consumed for each scheduled operation.

The Operation report is saved in the OperationsSchedule.csv file in the location that you specify when you create the publish profile.

Resource Schedule Report

The Resource Schedule produces a report which includes machine and crew information that is published according to the time fence that is selected.

The Resource Schedule is saved in the ResourceSchedule.csv file in the location that you specify when you create the publish profile.

Demand Fulfillment Report

The Demand Fulfillment report publishes detailed pegging information for each Demand line item. This includes the name of the customer, quantity of the desired item, available and requested dates.

The Demand Fulfillment report is saved in the Demand Pegging.csv file in the location that you specify when you create the publish profile.

Key Performance Indicators Report

Production Scheduling can publish KPIs to a .csv file, using the File|Publish tool. When you export KPIs, information about KPIs, including fill rate, stock outs, inventory turns, substitutions, resource utilization and throughput is exported.

The KPIs are saved in the KPI.csv file in the location that you specify when you create the publish profile.

HTML Dispatch Report

HTML file format generates an HTML Dispatch Report that can be published directly to the shop floor.

Shop floor personnel can use the dispatch list to organize production tasks for the day. The scope of the data that is published to the dispatch list is configurable and can be adjusted depending on the level of schedule data required on the shop floor. Once schedule information is dispatched to the shop floor, it is available to view on any networked workstation with a web browser.

The dispatch list index page is a centralized portal for shop floor personnel to view operation, item, work order, demand, and customer information for specific resources in order to plan daily production tasks. The dispatch list homepage is organized by machine, crew, and tool resources. Shop floor personnel can select the appropriate resource on the homepage and view detailed schedule information for that resource, depending on the scope of data determined prior to the publishing process.

The following schedule data can be published to the shop floor:

Category	*Data*
Operations	Operation ID, Description, Start Time, End Time
Items	Item, Item Description, Quantity
Work Orders	Work Order Number
Demand	Demand Number, Demand Due Date, Demand Item, Demand Quantity
Customer	Customer Name, Customer Number

 You can choose what headings you want in the report by selecting the check box in the profile's Data Selection window. Production Scheduling enables you to dispatch scheduling information to the shop floor efficiently without the need to write custom reports during the implementation phase. A web-based dispatch list can be published and deployed to a web server in your enterprise for shop floor personnel to use.

Work Order Schedule Report

Work Order Schedule report contains all the necessary information on work orders that you have scheduled.

The Work Order Schedule report is saved as an .xml file in the location that you specify when you create the publish profile.

Approving a Schedule

The schedule you want to publish must be approved before it is published.

To approve a schedule:

1. Access the Production Scheduling desktop.

2. Right-click on the scenario whose schedule you want to publish.

3. Select *Approve* from the available options.

    The green check mark beside the word *Approved* is set in a small box when the schedule is approved.

Publishing a Schedule Using Publish Profiles

To publish a schedule:

Note. Make sure that you have approved the schedule you want to publish

---

1. Access the Production Scheduling desktop.

2. In the Model Workspace, right-click the schedule that you want to publish.

3. Select *Approve* from the available options.

    Note. You cannot publish a schedule until it is approved.

    ---

4. From the File menu, select Publish.

5. In the Publish window, select the profile that you want to publish.

6. Click Publish.

Editing an Existing Publishing Profile

To edit an existing profile:

1. Access Production Scheduling desktop.

2. From the File menu, select Publish.

3. In the Publish window, select the profile that you want to edit.

4. Click the Edit button.

5. In the Edit Profile window, change the information you want to change.

6. Click Apply and then click OK.

Deleting an Existing Publishing Profile

To delete an existing publishing profile:

1. Access the Production Scheduling desktop.

2. From the File menu, select Publish.

3. In the Publish window, select the profile that you want to delete.

4. Click the Delete button.

5. Click OK when you are prompted to delete the profile.

Publishing a Schedule to Oracle E-Business Suite

You can publish a schedule to the planning server and to Oracle Applications Manufacturing modules using the Publish process. After the schedule has been solved, you can publish the results to the E-Business Suite. The schedule must be approved before it is published to Oracle E-Business Suite.

To publish a schedule to Oracle E-Business Suite:

1. Right-click on the schedule you want to approve for publishing.

2. Click Approve.

    The schedule name appears in bold in the Production Scheduling model workspace.

3. From the File menu, select Publish.

    The Publish window appears.

4. From the publishing profile list, select Release Schedule.

5. Click Publish.

    The schedule including the sequenced planned orders, jobs, scheduled resource, and specific changeovers are published to the planning server, and are automatically released to E-Business Suite.

Publishing Schedule to cMRO

The production schedule for maintenance organizations that use cMRO can be published using the Release Schedule publish profile. Users are able to selectively release work orders once they have been scheduled back into the cMRO system with updated dates reflective of the actual schedule and the firm status set by the user.

1. Right-click on the schedule you want to approve for publishing and select Approve.

2. Right-click on the schedule and select Edit Work Order Release.

3. Select the work orders you want to release.

4. Select the Publish icon from the toolbar or select Publish from the File menu.

5. Select Release Schedule from the Profiles list and click Publish. Only the select work orders are released.

Exporting Schedule Data

This section discusses how to:

• Set data export options.

• Export schedule data in XML format.

Setting Data Export Options

To set data export options:

1. Access the Export Schedule Options window.

2. Select the schedule data that you want to export from the available export options.

3. Click the Interval button to specify the time period or bucket, if necessary.

4. In the Interval Definition window, specify the following information:

Start Date	Specify the start date for the data. The default date is the beginning of the horizon.
Bucket	Select *Shift, Daily, Weekly,* or *Monthly* time buckets when exporting the current schedule.
End	Select the end of the interval. Select *No end date* for the end of the horizon. Select *End after* to select the number of time bucket occurrences that the exported schedule data represents. Select *End by* to specify a specific end date for the schedule data that you want to export.

5. Click OK to save the interval definition.

6. Click OK to save the data export options.

Exporting Schedule Data in XML Format

To export schedule data in XML format:

1. Access the Export Schedule window.

2. Select the XML option.

3. Select one or both of the following options:

Export Schedule	Export all schedule data to an XML file.
Export Purchase Order Recommendations	Select to export purchase order recommendations to an XML file.

4. Click OK.

5. In the Save As dialog box, enter the output XML file descriptive name or browse to an existing XML file.

6. Click Save.

Chapter 12

Integrating With Oracle Value Chain Planning Suite

This chapter provides an overview of Production Scheduling integration with Oracle Value Chain Planning Suite and discusses how to:

- Integrate Production Scheduling with the Oracle Value Chain Planning Suite.

- Import data to Production Scheduling.

- Use batch commands.

Understanding Production Scheduling Integration with Oracle Value Chain Planning Suite

Production Scheduling integrates with the Oracle Value Planning Suite (VCP) and Oracle ERP. This diagram illustrates how the application integrations with Oracle eBusiness Suite:

Note. Oracle ERP refers to Oracle E-Business Suite and Oracle JD Edwards EnterpriseOne.
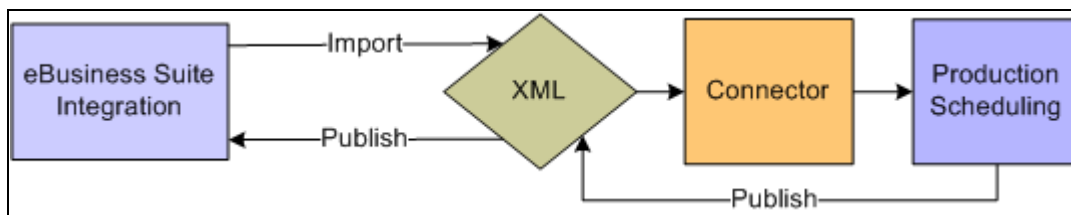
Production Scheduling integration with Oracle VCP

Integrating Production Scheduling with the Oracle Value Chain Planning Suite

This section provides an overview of Production Scheduling in the Oracle Value Chain Planning Suite (VCP) and discusses how to use Production Scheduling in the Oracle Value Chain Planning Suite.

Note. Oracle JD Edwards EnterpriseOne customers, please consult the *Value Chain Planning to JD Edwards Process Integration Pack User's Guide* for additional comments.

Understanding Production Scheduling in the Oracle Value Chain Planning Suite

Production Scheduling integrates with the Oracle Value Chain Planning Suite to enhance the capabilities of Oracle Value Chain Planning (VCP).

Production Scheduling uses a constraint-based approach to automated scheduling, to produce optimized production schedules. Unlike traditional automated scheduling tools that are limited to simple dispatch rules and have known bottlenecks, constraints in Production Scheduling can be assigned to every element such as resources, operations, and due dates in a schedule. By employing an advanced solver technology, feasible solutions can be found, if they exist, for virtually any floating constraint.

Business Processes

Through the seamless integration of Production Scheduling with Oracle VCP, you can create high level supply chain plans and refine these plans to create optimal production schedules that can be implemented using Oracle transaction systems.

Oracle ERP can send data to Production Scheduling one of two ways:

- From Oracle ERP to Production Scheduling.

- From Oracle ERP to ASCP to Production Scheduling.

   This allows Production Scheduling to receive planned orders from Advanced Supply Chain Planning (ASCP) when generating detailed production schedules.

When receiving data from ASCP, Production Scheduling receives planned orders. Once Production Scheduling creates a detailed shopfloor schedule, Production Scheduling sends the schedule directly to Oracle ERP. The results from the schedule can then be sent to ASCP to constrain the subsequent run of the ASCP plan.

See Integrating Production Scheduling, *Oracle Advanced Supply Chain Planning Implementation and User's Guide*

Implementation

The implementation of Production Scheduling in VCP can be divided into the following steps:

- Setting up the source.

- Setting up the destination.

- Setting up collections.

- Setting up profile options.

- Setting up plan options for Production Scheduling.

Using Production Scheduling in the Oracle Value Chain Planning Suite

You can use Production Scheduling in the Oracle Value Chain Planning Suite to take orders or planned orders from Oracle systems and create detailed, down-to-the-second manufacturing shop floor schedules and sequences. The Production Scheduling scheduling engine automatically identifies and resolves floating bottleneck problems to help maximize shop-floor efficiency and throughput while minimizing cost. You can review the resulting work orders and release them to Oracle transaction applications for implementation, and use them as constraints in subsequent ASCP planning. While ASCP creates plans for entire supply chains, Production Scheduling creates production schedules for individual plants.

To create production schedules, Production Scheduling requires an accurate representation of how your manufacturing plant works, including machines, crews, materials handling and other resource constraints. Because Production Scheduling is integrated with the Oracle Value Chain Planning Suite, you can obtain the required data from Oracle applications by:

1. Running collections to transfer enterprise and transactional data from Oracle transaction applications to the Operational Data Store (ODS).

2. Launching a Production Scheduling plan from Oracle eBS to transfer enterprise and transactional data from the ODS to create a model of your manufacturing plant in Production Scheduling and, if available, ASCP planned orders or Demand Planning demand forecasts from the Planning Data Store (PDS).

   When you launch a Production Scheduling plan, data is then imported into Production Scheduling on the client machine to create a manufacturing plant model.

   If ASCP is used in planning, Production Scheduling uses planned sales orders from ASCP as demand in the model. If ASCP is not used, Production Scheduling uses Demand Planning forecasts, if available, and sales orders from the transaction systems in the schedule horizon as demand.

   Note. When you launch a Production Scheduling plan from Oracle eBS, Production Scheduling is automatically installed on the client if it is not already installed.

After transferring data to Production Scheduling, you can create optimal production schedules for meeting the demand, review the schedules, and revise the schedules as desired.

You can then publish revised work orders from Production Scheduling to be released to Oracle transaction systems through the Planning Data Store (PDS).

The following diagram shows the integration processes of transferring data from Oracle transaction applications to Production Scheduling and transferring planned order from Production Scheduling to Oracle transaction applications:
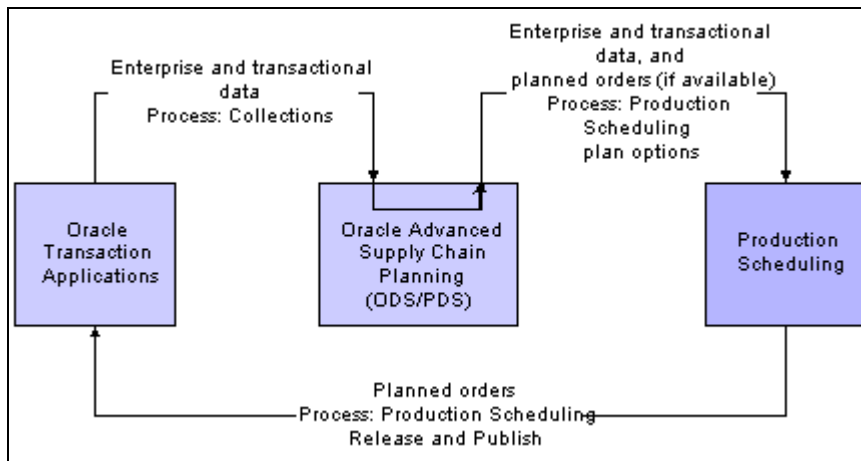
Oracle application integration with Production Scheduling

Importing Data to Production Scheduling

To import data into Production Scheduling, you can create an XML model.

You can create an XML model by using an ETL tool of your choice to create a series of CSV files. The CSV files represent the various objects in the Production Scheduling data model, such as crews, machines, tools, operations, routings, and so forth.

Using Batch Commands

You can use batch commands to publish data and perform a solve.

To solve and publish in batch mode.

1. Access a command prompt or a batch file.

2. Enter the following command:

    [directory path]\scp*[version number]*\ps\bin\PsBatch *pathfilename*

    where *[directory path]* is the path to where SCP is installed, *[version number]* is the is the version number of the SCP application, and *path filename* is the path and name of the batch input you have created to be used as input into the process.

    For example:

    scp*[version number]*\ps\bin\>PSBatch c:/scp/*[version number]*/ps/sample_data/samplebatchcommands.xml

    There is a sample batch command file entitled MyBatchScenarios.xml in the *installation drive and path* \scp*[version number]*\PS\sample_data directory.

3. Create your batch file, as displayed in this table:

Or you can refer to the example command file entitled SampleBatchCommands.xml in the *installation drive and path* \scp*[version number]*\PS\sample_data directory for the appropriate structure and nomenclature.

Batch file information is displayed in the table below.

Line Number	*Input*	*Description and Values*
1	`?xml version="1.0" encoding="utf-8"`	SCP supports .xml version 1.0 and utf-8.
2	`PSBatch version=`	Enter the version of Production Scheduling that you are using.
3	`model`	The opening tag for the model information.
4	clientDeploymentDir	Specifies the target folder. Note that the client data store is currently a hard-coded subdirectory of the client SCP installation folder.
5	outputFileName	This is the name of the location of the output file. You must specify a path and name of the .dxt file if you want to create the file. For example: C:\scp\ *[version number]*\PS\sample data\Bikes 12.1.1.dxt
6	`logFile`	Server log file location (not required as the required logs will also be contained within the outputFile).
7	`scenarioList`	This line begins the specifications for the scenarios that you want to publish. You can publish more than one scenario in a single batch file.
8	`scenario`	The opening tag for the model information.
9	`scenarioName`	Optionally, you can enter a name for the scenario. If you leave this field blank, Production Scheduling will automatically assign a name to the scenario.
10	`solveSchedule`	Directs whether the specified schedule should be solved.

Line Number	Input	Description and Values
11	inputFile	Enter the directory C:\scp*[version number]*\PS\sample_data\Bikesath and name of your file. The file must be in XML format.
12	kpiFile	Specifies the output location for the XML KPI report on the specified scenario.
13	/scenario	Ends the specifications for publishing the scenario. To publish more than one scenario, repeat the same information for the next scenario, beginning at line 8.

```xml
1 <?xml version="1.0" encoding="utf-8"?>
2 <PSBatch version="12.1.2">
3     <model>
4         <clientDeploymentDir>c:\scp\data_store\1138</clientDeploymentDir>
5         <outputFile>/user/serverDataStore/planDataStore/EbiPsSolved.dxt</outputFile>
6         <logFile>/user/serverDataStore/planDataStore/PSLog.txt</logFile>
7         <scenarioList>
8             <scenario>
9                 <scenarioName>Coffee</scenarioName>
10                <solveSchedule>yes</solveSchedule>
11                <inputFile>/user/serverDataStore/planDataStore/EbiPsModel.ops</inputFile>
12                <kpiFile>/user/serverDataStore/planDataStore/PsKpis.xml</kpiFile>
13            </scenario>
14        </scenarioList>
15    </model>
16 </PSBatch>
```

Example of a Production Scheduling batch file

Chapter 13

Scheduling Maintenance Work Orders

This chapter provides an overview of scheduling maintenance work orders in Production Scheduling and discusses how to:

- Schedule maintenance work orders taking into account visit details, milestones, and work breakdown structure.

- Support of complex task hierarchy relationships.

- Represent of the Work Breakdown Structure and assignment to work orders.

- Understand routings with maintenance work orders.

- Release Work Orders from Production Scheduling to cMRO.

Scheduling Maintenance Work Orders Taking into Account Visit Details, Milestones, and Work Breakdown Structure

Work orders from cMRO can now be imported and scheduled using Oracle Production Scheduling. New tables (6), housing such information such as visits, work breakdown structure, work order relationships, milestones have been added into the Operational Data Store (ODS) as this information is required to properly schedule cMRO work orders. Currently, this information is not available to standard planning collections from cMRO. An API into these new tables to augment collected data has been added to complement the existing collections framework as Production Scheduling requires this information to generate a maintenance schedule.

The overall architecture is represented in the diagram below.
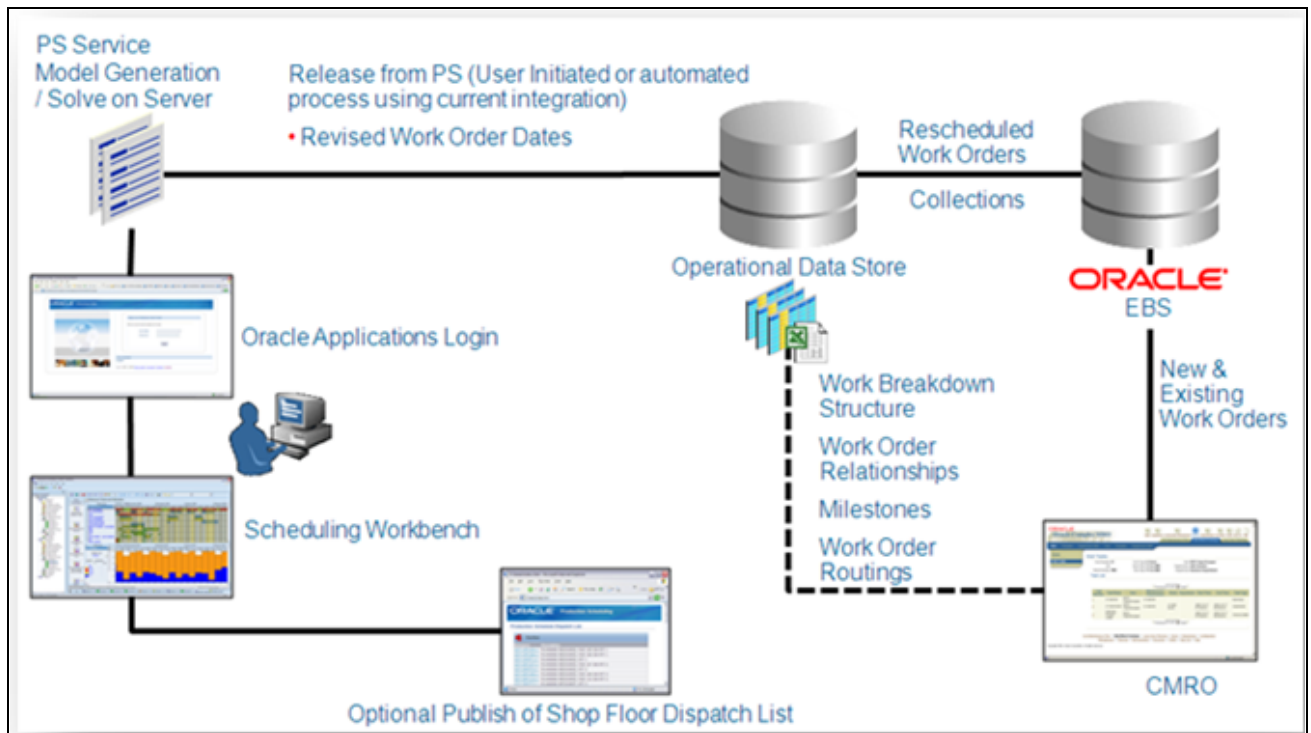
Production Scheduling integration with Oracle cMRO

For more information on cMRO, refer to the following:

• *Oracle Complex Maintenance, Repair, and Overhaul Implementation Guide* .

• *Oracle Complex Maintenance, Repair, and Overhaul User's Guide.*

Supporting of Complex Task Hierarchy Relationships

Visits have been added to Oracle Production Scheduling to allow for the scheduling, tracking, and grouping of tasks in the form of work orders and milestones. Visits and milestones are presented in the various views in Production Scheduling and allow the planner to leverage the rich set of existing functionality to review and refine the schedule. Visits and milestones also present new constraints which govern the production schedule that is generated.
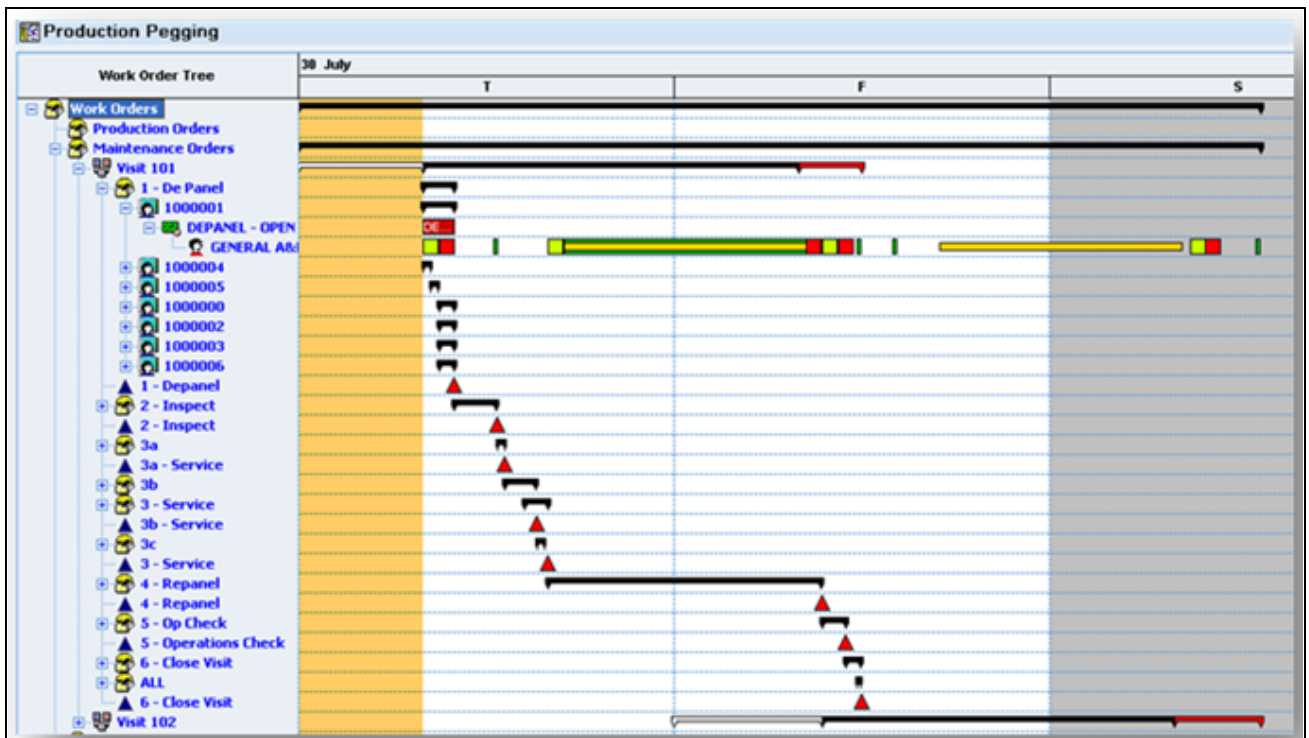
Example of complex hierarchy relationship

Task hierarchies are also introduced in the form of work order relationships where standard precedence constraints inherent in Production Scheduling can be leveraged to define the relationship between specific tasks.

In order to efficiently review the schedule, work order filters can be used and allow the planner to "slice and dice" the data in the various Gantt, Utilization, and Work Order views. This feature allows the planner to focus on specific tasks and answer questions such as "show me all work orders in stage 3a from July 2 - July 4 utilizing the Advanced Airframe and Power crew which are firm". While this feature has been added to facilitate cMRO workflows, it will benefit many of our existing manufacturing customers as they will also be able to filter on many of the existing constructs already available in work orders.

Once the cMRO model is reviewed and analyzed in Production Scheduling, work order dates generated by the schedule and can than be released back to cMRO for execution.

Representing of the Work Breakdown Structure and Assignment to Work Orders

The logical grouping of work can be accomplished by leveraging the cMRO Work Breakdown Structure, which is passed into Production Scheduling. The Display Parameters tab in the Plan options screen is where this can be used to create logical groupings of work orders. The Selected attributes are used to group work orders in folders in the work order editor in production scheduling and allows the planner to group the work orders based on the Work Breakdown Structure.

The work breakdown structure is populated as work order attributes and can be viewed in the work order editor. This allows for advanced filtering in Work Order Filters where they can be used to group work orders in the various views when analyzing the schedule.

The solution may be applicable for other industries with similar requirements, but focus has been on solving aircraft short term maintenance scheduling problems.

See Integrating Production Scheduling, *Oracle Advanced Supply Chain Planning Implementation and User's Guide*

Understanding Routings with Maintenance Work Orders

When working with maintenance work orders, routings in Production Scheduling are created directly from the work orders themselves. All routings with be stored in a Work Order Routings folder under Routings.

Releasing Work Orders from Production Scheduling to cMRO

Production Scheduling allows you to selectively release work orders once they have been scheduled back into the cMRO system with updated dates reflected by the actual schedule and the firm status you defined. This is performed by approving the schedule, publishing the work orders using the profile setting defined for your cMRO process. The revised Work Order Dates are visible back in cMRO.

Index

A

alerts
    expanding messages 199
    filtering messages 200
    grid 199
    navigate to demand 200
    tree
        model workspace 198
    understanding 198
    work with 197
all of sets
    creating 67
alternate
    resources in Production Pegging 191
assembly buffers 80
availability
    defining 82

B

batchable
    resources 35
batch commands
    using
        Oracle VCP 242
buffers
    assembly 80
    assigning 78, 82
    assigning in routing diagrams 81
    assigning to routings 81
    capacity constrained resource 79
    defining attributes 81
    shipping 79
    supplier 80
    validation rules 80
business processes
    interactive events 4
    non-interactive events 3
    of Production Scheduling 2

C

Calendar editor
    configuring model properties 82
    creating a global 83
    data editors 19
calendar events
    adding 83
    deleting 85
    modifying 84
calendar horizons 83
calendars
    adding events 83
    adding resource events 223
    creating a global 83
    deleting events 85

    modifying events 84
campaign run optimization
    model considerations 157
    setting up 133
    understanding 151
Campaign Run Optimization 151
    changeover costs 152
    changeovers between cycles 154
    changeover time 152
    example 152
    ideal run sequence 155
    inventory carrying costs 152
    inventory stockout costs 152
    minimum cycle time 154
    minimum run length specification 159
    multi-stage campaign run optimization
        limitations 160
    safety stock violation costs 152
    setting up 160
    user defined ideal run sequence 155
capacity
    defining for resources 38
capacity constrained resource
    buffers 79
Changeover editor
    data editor 19
changeover rules
    configuring model properties 121
changeovers
    between cycles 154
    CRO and costs 152
    defining rules 121
    filter changeover rules 123
    find matching changeover rules 122
    re-ordering rules 122
    time 152
    viewing 167
    viewing in schedule views 167
cMRO
    publishing a schedule to cMRO 236
combination view
    adding 195
combined
    views 182
combined views
    viewing 196
comparing
    schedules 205
configuring
    desktop components 28
    diagrams 24
    operation and routing diagram layout 126
    Production Pegging layout 28
    toolbars 27
constraints
    defining for items 48
    precedence 71
    reactivating for resources 189
    relaxing for resources 189
cost indicators
    key performance indicators 203
costs
    defining for operations 64

R