



CHAPTER 8

Planning Fundamentals

The Oracle Applications suite currently offers two products for material and capacity planning in a manufacturing or distribution environment: Master Scheduling/MRP and Advanced Supply Chain Planning. Master Scheduling/MRP offers single-organization, unconstrained planning of material and resources. We'll use the term *MRP* to refer to this planning tool. Advanced Supply Chain Planning (ASCP) is the next-generation planning tool; it offers multi-organization planning and the option of constraint-based and optimized plans.

This chapter describes the basic logic common to both planning methods and details the common setup needed by both methods. Subsequent chapters describe the unique characteristics of MRP and ASCP.

NOTE

In past releases, Oracle sold a product called Supply Chain Planning, which provided the capability of planning across a supply chain, but without the constraints, optimization, and variable granularity of the Advanced Supply Chain product. With Release 11i, this is no longer sold as a separate product, so it will not be explicitly covered in this book. However, much of material describing Advanced Supply Chain planning applies to the older supply chain planning; notable exceptions are the constraint and optimization capabilities of ASCP and the treatment of sourcing rules, which are described in Chapter 11.

Basic Planning Logic

The basic logic of both MRP and ASCP is time tested and well documented in numerous sources. Though not a comprehensive discussion of basic MRP logic, what follows is a brief overview of the process. The process is illustrated in Figure 8-1. It involves the recognition of demand or requirements, the netting of those requirements against available and scheduled quantities, and the generation of recommendations to meet those requirements. In a manufacturing environment, the process proceeds top-down through the bill of material, so that recommendations at one level form the basis of the requirements on the next level of components, and so on.

Recognition of Demand

Demand initially starts as a statement of what a company expects to sell or ship. This type of demand represents what you know or think you are likely to sell outside

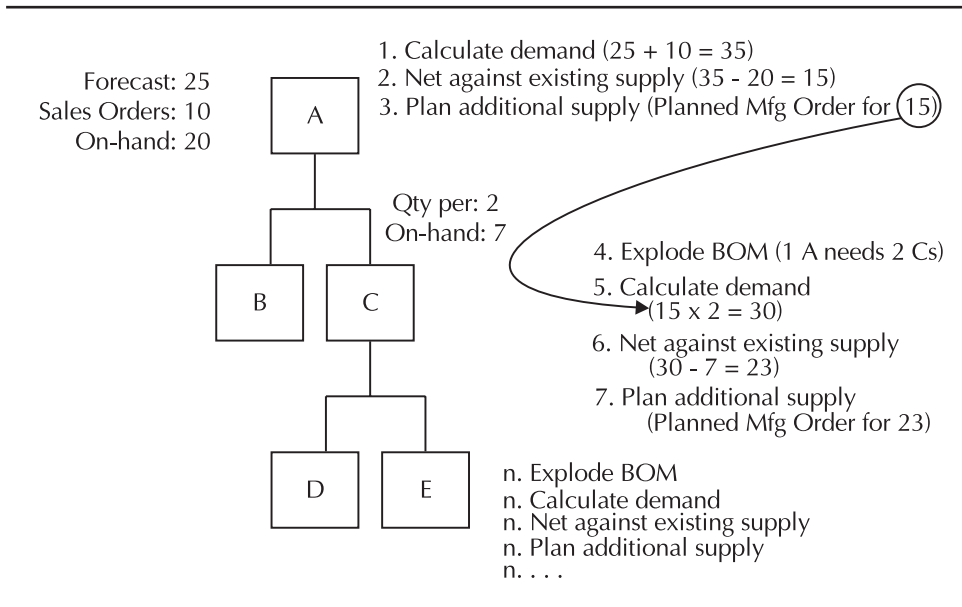


FIGURE 8-1. Basic netting, planning, and explosion logic

your company and is traditionally called *independent demand* because it is to some extent outside of your control. In a manufacturing business, the independent demand for your finished goods determines the demand on the component parts; this is called *dependent demand*, since it depends on, and can be calculated from, your bills of material.

Independent demand is typically some combination of forecast and actual (Sales Order) demand. Depending on the stocking strategy you employ for a particular product, the mix of forecast and actual demand will vary; in a make-to-stock strategy, you will rely heavily, perhaps exclusively, on forecast demand. In a make-to-order strategy, you will rely heavily on actual sales orders. Furthermore, the ratio of forecast to actual will likely vary over the time horizon, as shown in Figure 8-2—in the near term, you might rely most heavily on actual demand, while using forecast demand at the mid- and long-range points of the time line.

Oracle Applications allow you to have alternate forecast scenarios, so the forecasts you want to use to plan, as well as the actual demand you want to recognize, are stated in a Master Demand Schedule (MDS). The MDS is discussed in more detail, in the “Master Scheduling” section.

Dependent demand is demand calculated from the explosion of your bill of material. An important premise of MRP planning is that you should calculate, not forecast, dependent demand. For example, one automobile requires four tires; if you forecast or sell one auto, you know you will need four tires to build it. There’s

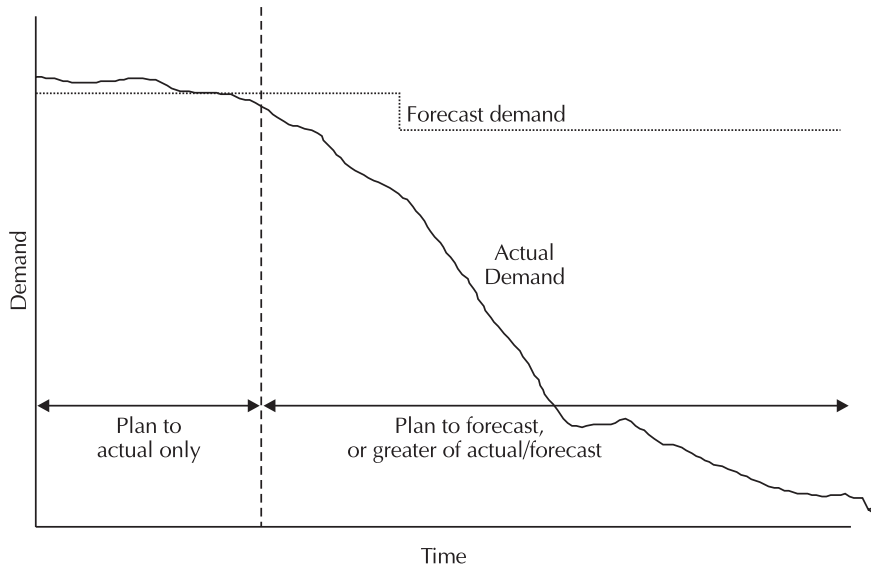


FIGURE 8-2. *Actual and forecast demand*

no point in forecasting tires; you can calculate the requirements from the bill of material. Because they represent an attempt to predict the future, forecasts are notoriously inaccurate; forecasting what you can calculate can only compound problems.

It is possible that one item can have both independent and dependent demand. If you were to sell tires as service parts, in addition to including them on a manufactured auto, you would want to forecast the independent *portion* of the demand on tires while calculating the dependent demand.

Gross to Net Calculation

Once the demand is known, the next step is to determine how much of that demand can be satisfied from existing stock or existing orders. It's simple arithmetic; if you have gross requirements (total demand) for 100 units of a product, and you have 25 on hand, the net requirements are 75.

Time-phasing makes this arithmetic a little more interesting. If that same demand were spread out over four weeks (25 per week), your 25 on-hand would satisfy the first week's demand. If you had an existing order (e.g., a manufacturing order) to build 25 more in week 3, the system would expect you to reschedule that order to satisfy the demand in week 2 and would then see net requirements for 25 each in weeks 3 and 4.

Recommendations

In the scenario described earlier, the planning process would recommend that you reschedule the existing order from week 3 to week 2 and would recommend that you create new orders to satisfy the demand in weeks 3 and 4.

In this process, planning will respect any order modifiers you've attached to the item, such as minimum or maximum lot sizes. It will look at any anticipated shrinkage factors you've associated with the product and inflate the planned orders to compensate. And it will respect time fences to allow you to stabilize your production or procurement schedules in the short term. It is the application of all these factors that makes planning beyond the scope of paper and pencil for all but the simplest scenarios. Order modifiers, shrinkage, and time fences are described later in this chapter.

Bill of Material Explosion

In a manufacturing environment, once the top-level demands are known, the component demands are calculated based on the explosion of the bill of material. The component demand (plus any independent demand for the component) constitutes the gross requirements at the next level. The planning process then continues to calculate net requirements and plan orders at each level in the bill of material.

Planning Cycles

You can use Oracle's planning products in numerous ways to generate distribution plans, master schedules, and manufacturing (MRP) plans for your enterprise. The possibilities are almost endless, but we'll discuss two typical examples here.

Figure 8-3 shows the traditional planning cycle. In this cycle, you create forecasts and book sales orders, load the selected forecasts and orders into a Master Demand Schedule, use the MDS as input to planning a Master Production Schedule, and submit that MPS to MRP. This cycle separates master scheduling and MRP so that you can firm up a master schedule before submitting it to MRP. You must synchronize these plans by carefully considering the timing of the running of each plan.

Figure 8-4 shows a *holistic* planning cycle; in this cycle there is only one plan that satisfies the needs of distribution planning, master scheduling, and manufacturing. You still generate forecasts, book orders, and load demand schedules, but you use one or more demand schedules from multiple organizations to generate one complete supply chain plan.

As noted earlier, numerous variants of these cycles are possible. A company that does not need to separate the Master Production Schedule from the MRP Plan can run MRP directly from the MDS. Within the supply chain plan, master schedules can be used to introduce some stability and local control into the plan. But all plans require some initial statement of demand, usually a combination of forecast and

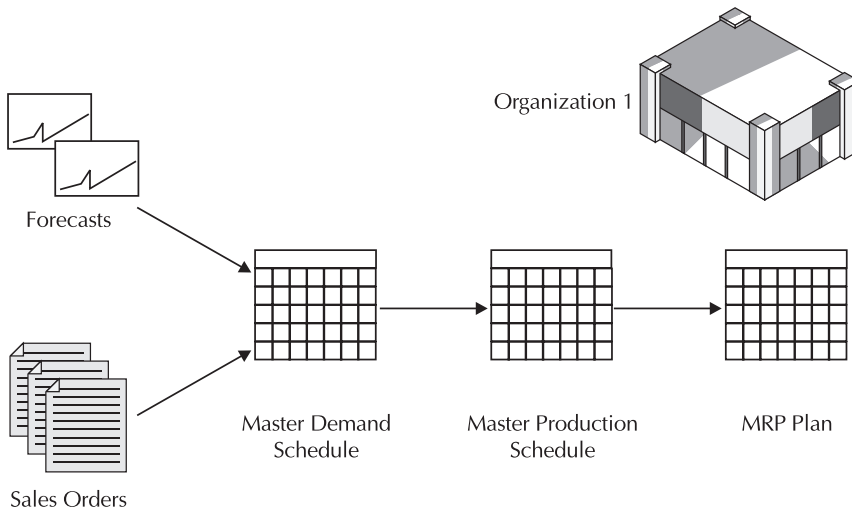


FIGURE 8-3. *Traditional planning*

sales orders. All plans require some sort of master schedule to identify the demand that drives the plan. They will plan for effectivity date changes in your bills of material and plan for both discrete and repetitive production. These common elements are described in the following pages.

Forecasting

For most businesses, forecasts are necessary at some point in the time horizon. Even a company with a pure make-to-order strategy (if such a company exists) will probably want to use forecasts of future demand for long-range planning. Forecasting is at least as much art as science, and the art of forecasting is beyond the scope of this book. But we will describe how forecasts are represented, generated, and used within the manufacturing and supply chain planning process.

Oracle Demand Planning (described in Chapter 9) generates forecast data. In addition, the Oracle Inventory and Master Scheduling/MRP applications provide basic methods to generate forecasts from historical data, described in the section titled “Generating Forecasts from Historical Information.” Forecasts can also be entered manually; imported through an open interface, described in the “Open Forecast Interface” section; or loaded via an API, described in the section titled “Forecast Entries API.”

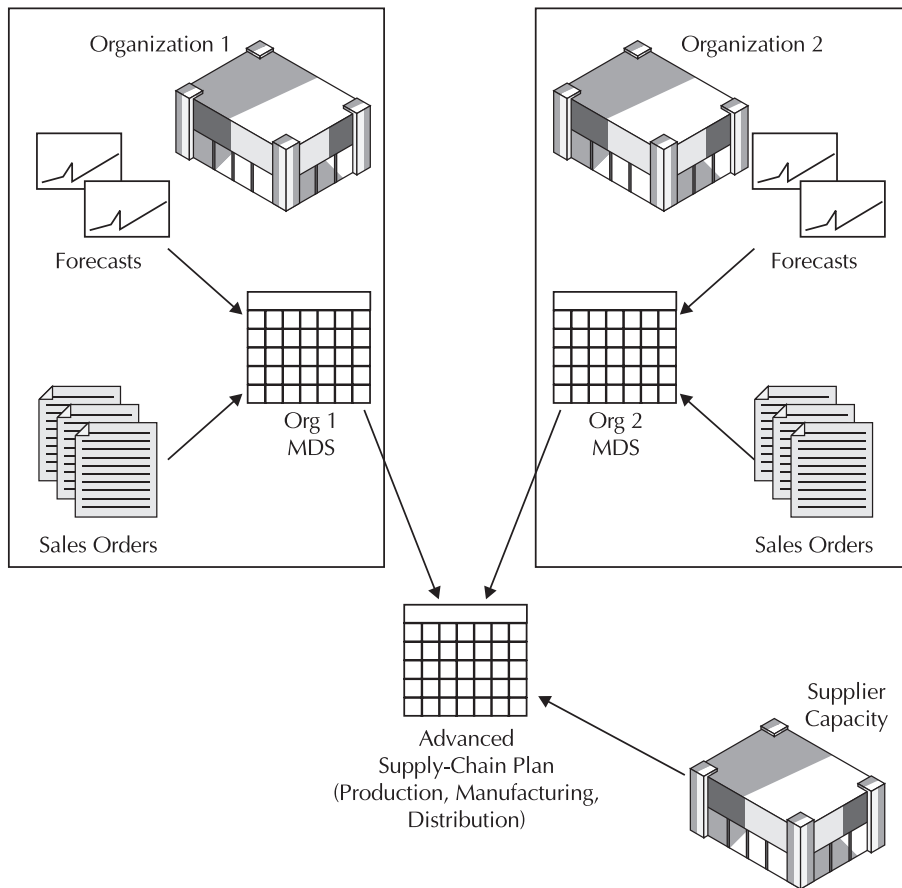


FIGURE 8-4. *Holistic planning*

Forecast Terminology and Structure

To understand the use of forecasts in planning, it is helpful to understand the terminology and data structure Oracle uses to represent forecasts. One *forecast* typically contains multiple items; each item has multiple *entries* that represent a dated demand for that item. For ease of use and to control forecast consumption, forecasts are grouped into *forecast sets*.

Forecasts, Forecast Items, and Forecast Entries

Within Oracle applications, forecasts are identified by unique names. For example, you might have separate East, West, and Central forecasts, or you might have different forecast names by product line or distribution channel. You might have Optimistic and Pessimistic forecasts or simulations of various scenarios. Each such forecast is identified by a name.



NOTE

Give some thought to the conventions you will adopt for naming forecasts. Unlike many names in Oracle applications, forecast names (along with master schedule and plan names) are limited to 10 characters in length and cannot be changed.

A forecast will typically contain multiple item numbers, called forecast items, and each item will have multiple forecast entries, each representing the forecast demand for a particular day, week, or monthly period. (Oracle seems reluctant to use the term *month* unless it refers to an actual calendar month; in the context of planning, you might be using a 4-4-5 or 5-4-4 calendar, whose periods don't line up with actual calendar months.)

Your workday calendar controls the valid dates for forecast entries. Daily forecasts can be entered for working days only; weekly or period forecasts will be dated on the first day of the week or period.



CAUTION

If your workday calendar specifies a Quarterly Type of "Calendar months," Oracle does not currently support the use of a "Period" bucket type for forecasts.

Forecast Sets

Oracle requires you to group forecasts into forecast sets, which are simply collections of forecasts. Forecast sets serve as both a convenience (you can refer to multiple forecasts by referencing their set) and a method to control forecast consumption, described in the following. This structure—forecast sets, forecast names, forecast items, and forecast entries—is illustrated in Figure 8-5.



NOTE

Forecast set names are defined in the same database table that defines forecast names, so they share the same size restrictions. In addition, forecast set names must be different from forecast names.

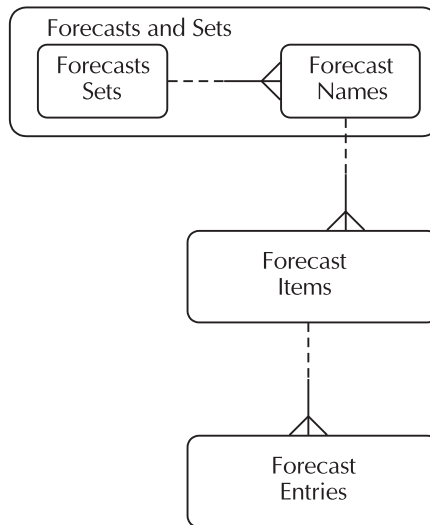


FIGURE 8-5. *Forecast data model*

Forecasts are always organization/item-specific; they represent anticipated demand for an item from an individual organization. But they may be defined at a more detailed level, if necessary. You can define a forecast at the customer level; such a forecast represents demand for an item from an organization to satisfy the orders of a particular customer. Or, you can define the forecast at the customer address level, where demand is maintained by customer ship-to address.

At the forecast set level, you define the default bucket size (day, week, or period) for the forecasts in the set, though this can be easily overridden for any forecast entry. You define the level (item, customer, or address) for which all forecasts in the set will be maintained. And you define certain controls that affect forecast consumption; these will be discussed later, when we discuss the forecast consumption process.

At the forecast level, you define the level identifiers (e.g., the customer name if the forecast is maintained at that level) and the optional demand class that will further control forecast consumption.

Generating Forecasts from Historical Information

The Oracle Master Scheduling/MRP and Inventory Applications provide two basic methods to generate forecasts from historical data: statistical forecasting and focus forecasting. While both methods use proven algorithms, they are basic in the level of control they provide. Unlike Oracle Demand Planning, they are based strictly on

the history of issue transactions (not order bookings); therefore, if you have shipped late in the past, these methods will forecast how you can ship just as late in the future! And they offer none of the outlier detection or collaborative features of Oracle Demand Planning. Nevertheless, they may meet simple requirements.

Statistical Forecasts

Oracle provides basic statistical forecasting capability as part of the Inventory and Master Scheduling/MRP applications. Statistical forecasts can span multiple periods and can recognize trend and seasonality.

Focus Forecasts

Focus forecasting provides a “sanity check” on the current month’s forecast. Focus forecasting examines five different forecast models against past history, determines the model that would best have predicted the history, and uses that model to generate a forecast for the current period. (The process generates a “forecast” for multiple periods, but they are simply copies of the current period forecast.)


Generating Forecasts

You generate a forecast by first defining a forecast rule that specifies a bucket type, a forecast method (statistical or focus), and the sources of demand. For the statistical forecast method, the forecast rule also specifies the exponential smoothing (alpha) factor, trend smoother (beta) factor, and seasonal smoothing (gamma) factor. Then, in the Generate Forecast window, you specify the following:

- Name of the forecast you want to populate.
- Forecast rule that determines the bucket, forecast method, and demand stream.
- Selection criteria to identify the items for which you want to generate a forecast. You can select all items, all items in a designated category set, items in a specific category within a category set, or an individual item.
- An overwrite option, which determines what happens to existing forecast information. All Entries deletes all existing forecasts before repopulating the forecast data. Same Source Only deletes only forecast entries that were

previously generated from the same source that you are about to load. None does not delete any existing entries.

- Start date and cutoff date for the forecast information that you are about to generate.



NOTE

Choose the Overwrite option carefully; if you do not explicitly delete existing information through your choice of the Overwrite option, you might duplicate information already in your forecast.

Open Forecast Interface

To populate forecasts using the Open Forecast Interface, you must first define your items, organizations, forecast sets, and forecast names. Then, load your forecast data into the table MRP_FORECAST_INTERFACE. Required data includes

- INVENTORY_ITEM_ID
- FORECAST_DESIGNATOR
- ORGANIZATION_ID
- FORECAST_DATE
- QUANTITY
- PROCESS_STATUS

Set the PROCESS_STATUS to 2 to designate data waiting to be processed.

The Planning Manager (described later in this chapter) periodically checks the MRP_FORECAST_INTERFACE table to determine if there is any new data to be processed. If the Planning Manager detects new data, it launches the Forecast Interface Load Program. This program validates the data in the interface table and loads valid data into the specified forecast. If the load program detects an error, it sets the PROCESS_STATUS to 4 and enters a text message in the ERROR_MESSAGE column in the interface table. Use SQL*Plus, or a custom program, to view and correct invalid data in the interface.

Forecast Entries API

You can use the forecast entries API to integrate with other software systems. This section describes the major steps in using this API.

1. Construct the following PL/SQL table for all the information that is associated with the forecast:

Variable Name	Type	Meaning
p_forecast_entries	A PL/SQL table that contains records of type mrp_forecast_interface_pk. mrp_forecast_interface	These are the actual forecast entries.
p_forecast_designators	A PL/SQL table that contains records of type mrp_forecast_interface_pk. rec_forecast_desg	These are the forecast designator information.

2. Call the function `mrp_forecast_interface`, with the parameters that you constructed in step 1. You don't have to supply both the parameters in every case. For example, if you are inserting new forecast entries or updating existing forecast entries, you would provide all the information. On the other hand, if you are deleting all the forecast entries, you need only provide the forecast designator.

```
mrp_forecast_interface_pk.mrp_forecast_interface
(p_forecast_entries
-- INOUT PL/SQL Table for forecast entries
,p_forecast_designators
-- IN PL/SQL Table for forecast designators
);
```

3. The function returns true if the API completed successfully and false if there were failures.

Forecast Consumption

In order to avoid overstating demand as actual orders are booked, forecasts are typically reduced by the amount of the sales order demand. The process is called forecast consumption. Then, when actual orders are combined with forecasts in a master schedule, the demand statement remains accurate. *Forecast consumption* also serves as a measurement of the forecast's accuracy—if you've sold less than you forecast, the forecast will not be fully consumed; if you've sold too much, the result will be reflected as overconsumption at the forecast set level.

Forecast consumption is one of the functions of the Planning Manager, described later in this chapter. At its simplest, forecast consumption reduces a forecast demand by the amount of the actual sales order demand in that period. The remaining forecast quantity is referred to as the current forecast; the original forecast quantity remains visible for comparison and analysis.

The first rule of forecasting is that the forecast will be wrong. It is highly unlikely that a customer order will exactly match your forecast quantities (although supply chain collaboration might help ameliorate this problem). So, to make the process yield usable results, a number of factors influence the supposedly simple process of forecast consumption.

Forward/Backward Days

To compensate for the discrepancy between customer orders and forecasts, you typically provide a “window” within which the forecast will be consumed. This window is expressed as a number of working days, as determined by the shop calendar. Within this window, forecast consumption will first look backward and then forward to match the forecast with the actual demand. You define this window as an attribute of each forecast set.

Consider this example:

Days	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Forecast	5	5	5	5	5			5	5	5	5	5			5	5	5	5	5	
Demand										30										

= Non-work Day

With no window for consumption, the sales order demand on day 10 would consume the entire forecast of 5 on that day and then “have no where to go.” Thus, the total demand for the second week would be 50: 30 actual and 20 of the remaining, or unconsumed, forecast, as shown here:

Days	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Forecast	5	5	5	5	5			5	5	5	5	5			5	5	5	5	5	
Demand										30										
Unconsumed	5	5	5	5	5			5	5	0	5	5			5	5	5	5	5	

Daily buckets; Backward Days = 0; Forward Days = 0 = Non-work Day

With forward and backward days set to 3, the actual demand would “search” for forecast if there was not sufficient forecast quantity remaining on the demand date. The process always goes backward first (assuming that you are late in realizing your forecast), and then forward if necessary. So the demand of 30 on day 10 would consume 5 on day 10, then search backward consuming forecast on day 9, 8, and 5 (day 6 and 7 are non-work days). Since the backward consumption only accounted for 20 of the 30 demanded, the process would then search forward and consume forecast on day 11 and 12. The result is shown here:

Days	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Forecast	5	5	5	5	5			5	5	5	5	5			5	5	5	5	5	
Demand										30										
Unconsumed	5	5	5	5	0			0	0	0	0	0			5	5	5	5	5	

Daily buckets; Backward Days = 3; Forward Days = 3 = Non-work Day

Forecast Bucket Type

The bucket type of a given forecast entry allows an order anywhere within that bucket to consume forecast. If the forecasts in the previous example were expressed as weekly instead of daily forecasts, the order on day 10 would consume all the forecasts for the week starting with day 8, even with no backward/forward days:

Days	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Forecast	25							25							25					
Demand										30										
Unconsumed	25							0							25					

Weekly buckets; Backward Days = 0; Forward Days = 0 = Non-work Day

Note that the consumption window is always expressed in days, regardless of the forecast bucket type. Thus, in the example above, a backward parameter of 2 days would not change the consumption, since 2 days backward only moves the consumption window back to day 8. A backward parameter of 3 days, however, would move the consumption window back to day 5, which falls in the previous week. Since the forecast bucket is a week, that week’s forecast would then be consumed. Thus, a backward consumption parameter of 3 days (or more) would result in the following:

Days	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Forecast	25							25							25					
Demand									30											
Unconsumed	20							0							25					

Weekly buckets; Backward Days = 3; Forward Days = 0 = Non-work Day

Forecast Level

As mentioned earlier, you can forecast at three levels: item/organization, item/organization/customer, or item/organization/customer address. Sales Order demand only consumes the appropriate level; i.e., if your forecast is maintained at the customer level, only a sales order for the customer associated with the forecast will consume that forecast. Likewise, if you maintain a forecast at the customer address level, only a customer order with a matching ship-to (or bill-to) address will consume that forecast.

Outlier Update Percent

The dynamics of your particular market will determine how much of a forecast you will want a single sales order demand to consume. For example, if you sell to major retailers or distributors, you might expect a single order to consume most or all of a forecast. But if you are selling computers directly to individual consumers, it might be highly unusual for a single sales order to consume more than a small fraction of a week's forecast. In this situation, an unexpectedly large order would be called an outlier because it represents demand that falls outside of the expected pattern. If you were to allow such an outlier to consume the entire week's forecast, you might be understating demand—you would show no remaining forecast, even though you might expect that your normal demand patterns would persist, and additional orders would eventually be placed according to expectations.

To prevent such unwanted forecast consumption, Oracle lets you specify an outlier update percent; this determines the percentage of an individual forecast that a single demand (i.e., a single order or shipment line) can consume. For example, if you had an original forecast quantity of 100, but had set an outlier update percent of 10, no sales order line could consume more than 10 percent, or 10 units, of the original forecast quantity. A sales order for 20 units, therefore, would only consume 10 units of the forecast, leaving an unconsumed balance of 90 units for that forecast entry. Note, however, that the forward/backward days would still be applied; the excess amount of the order could still consume other forecast entries based on your consumption window.

Consumption Within Sets

A given demand can consume multiple forecasts, but it will consume forecast entries from only one forecast within a forecast set. For example, if you have an Optimistic, Realistic, and Pessimistic forecast, you should define them in different forecast sets; each forecast could therefore be consumed by your sales order demand so you could gauge the accuracy of each forecast. (You would, of course, use only one forecast for a given planning scenario, or you would be overstating demand.)

Demand Class

Forecast sets control forecast consumption because a given demand will consume (or attempt to consume) forecast entries in only one forecast within the set. Sometimes your forecast level will control which forecast is consumed within a set—if your forecasts are defined for each customer, the customer on the sales order will determine the forecast to consume. However, if you're forecasting at the item level only, and if you have the same item in multiple forecasts within a set, you will need to control which forecast should be consumed. You control forecast consumption with *demand classes*.

A demand class is simply a name you attach to a forecast and sales order demand; then sales order demands will only consume forecasts with a matching demand class. In the absence of a demand class, a sales order will consume the first forecast it finds (alphabetically) in a forecast set. Thus, if you had an Asia, Europe, and North America forecast in a forecast set, a sales order from New York would consume the Asian forecast, since Asia sorts alphabetically before North America. You can control the consumption by defining demand classes that correspond to each forecast—you might simply use the same name for the demand class and the matching forecast—and then apply the demand class to the appropriate sales order line.

Order Management lets you attach a demand class to an order type or to a customer address; defaulting rules would apply the desired demand class to a sales order line. If you forecast geographically, you would probably determine the demand class from the customer address; if you forecast by distribution channel, you might determine the demand class from the order type.

An order will normally consume only forecasts with a matching demand class; however, if there is no forecast with a matching demand class, the order will consume forecasts that have no demand class (i.e., the demand class is null). You might use this capability to maintain specific forecasts for some of your major sales channels, for example, while maintaining a generic forecast for all other demand.

Similarly, an order with no demand class will normally consume a forecast with a null demand class. If you prefer to have such orders consume from a specific demand class, you can specify the demand class as the Default Demand Class on the Organization Parameters form.

In some cases, there will be orders that you will not want to consume forecast. For example, you might have received an order that you know is due to production

problems at a competitor or due to some other event that you did not anticipate in your forecast. Such orders (sometimes called “bluebirds”) should not consume forecast; they represent excess or abnormal demand and should be added to the forecast. You can accomplish this with demand classes as well—define a demand class with a name such as Abnormal and assign it to the abnormal demands that you recognize. Create a dummy forecast for the Abnormal demand class. The Abnormal demand class and forecast will keep the demand from consuming another forecast; if you exclude the Abnormal forecast but include all sales orders (including the “abnormal” ones) in the master demand schedules driving your plans, you will properly recognize the abnormal demand in addition to your normal forecasts and orders.

Aggregate Forecasts

In some cases, it is easier to forecast at an aggregate level and use an anticipated distribution percentage to calculate the forecast for an individual item. Neptune Inc., for example, sells computers with a mix of processors, hard drives, CD drives, etc. Their sales and marketing staffs can forecast with some success the total number of computers they will sell, and they have a good idea of the relative mix of options. However, trying to forecast individual combinations of features is nearly impossible. In this case, the model and option bill of material that defines the options that customers can choose includes a distribution percentage on each component in the bill. This lets Neptune enter a forecast for the model item, and the forecast is exploded to calculate a forecast quantity for each option.

You can use a similar process even for items that you make to stock. You can group similar products in a product family or planning bill and enter an aggregate forecast for the family or top level of the planning bill. Both product families and planning bills let you aggregate forecast; additionally, product families can be used in the design of flow lines (discussed in Chapter 7) and to enable two-level master scheduling, described later in this chapter.

Forecast Control

Forecast control is an item attribute that controls the forecast explosion process. The attribute has three possible values:

- **Consume and Derive** This setting implies that the forecast can be calculated, or derived, from the explosion of an aggregate forecast; since the item will have a forecast, it is eligible to be consumed. Note that this setting does not *require* that forecast be derived, only that it is possible. It is a good choice as a default setting for your item templates, as it allows forecasting, consumption, and calculation from a model bill, planning bill, or product family.

- **Consume** A forecast control setting of Consume implies that forecast will be maintained separately for the item, but not derived from explosion of an aggregate forecast. This might be useful if you have an item that you want to forecast separately, but want to include on a planning or model bill for completeness.
- **None** None implies that no forecast is maintained for the item.



NOTE

To explode and derive a forecast through multiple levels of a planning bill or model and option bill, you must set the Forecast Control attribute for all intermediate levels in the bill to Consume and Derive. If you “break the chain,” forecast explosion will stop.

Forecast Explosion

Forecast Explosion is the process that takes an aggregate forecast and generates individual item forecasts for the components of a planning bill, model and option bill, or product family. You can initiate the process at several points: when you copy a forecast or when you load a forecast into a master schedule.

Forecast Explosion takes each forecast quantity at the top level and explodes it through the bill of material, multiplying by the Distribution Percent and the Quantity per, if any. It continues down the bill of material until it has created forecast for the first standard item it finds down each branch. In other words, once Forecast Explosion finds a standard, stockable component, the explosion of that branch of the structure is finished; it does not generate forecast for components of subassemblies, for example.

Master Scheduling

The planning process always starts with a master schedule. Like forecasts, each master schedule is identified by a 10-character name. Oracle supports two types of master schedules: Master Demand Schedules and Master Production Schedules.

While the distinction between a Master Demand Schedule and a Master Production Schedule is unfamiliar to many people, it is a useful concept. In many businesses, especially for seasonal products, it is not feasible to produce exactly to meet the customer demand. In such cases in particular, it is helpful to separate demand from supply.

Master Demand Schedules and Master Production Schedules

A Master Demand Schedule (MDS) is an anticipated shipment schedule; it is the statement of demand you want to recognize for a particular planning run. A Master Production Schedule (MPS) is a production plan, a statement of how you plan to schedule production. It might also be called a production forecast.


You define master schedules with several forms within the Oracle Applications suite. Use the Master Demand Schedules form to define an MDS; all that is required is a name and description. You can optionally associate a Demand Class with an MDS and indicate if shipments should automatically relieve the MDS. You would probably want to enable shipment relief for the master schedule that represents the demand you actually plan for, but you might choose to turn off relief for an MDS that you use for simulation or historical purposes. (Master Schedule relief is discussed in more detail later in this chapter, in the section titled “Schedule Relief.”)

The form you use to define an MPS depends on the type of plan you will execute. Use the Define MPS Names form for a master production schedule that you will use for single-organization, Master Scheduling/MRP plan. This form is found under the Material Planning menu seeded in Oracle Applications. If you’re running the older supply chain plan, you must use the Define Multi-Plant MPS Names form, found under the Supply Chain Planning menu. Like an MDS, you can associate an MPS with a demand class and enable or disable relief. In addition, for an MPS, you must specify if it is to be considered in the calculation of Available to Promise quantities in Order Management and Inventory and whether or not it is considered a Production plan. Designating an MPS as a Production plan allows automatic release of orders during the planning process.

To define an Advanced Supply Chain plan, use the MPS Names form under Advanced Supply Chain planning; this will be discussed in greater detail in Chapter 11.

You can enter master schedule quantities manually, but most often you will use concurrent programs within Oracle Applications to populate your master schedules with the appropriate data. An MDS will typically contain a combination of forecasts and sales orders. You can also copy one or more master demand schedules into a new MDS; this can be useful for simulations or for retaining historical data.

As a statement of production, an MPS will represent your production plans. In an environment where you can produce exactly to customer demand, you might simply copy your MPS from your MDS and then modify it manually. Like an MDS, you can also copy one or more master production schedules into a new MPS for simulation or to retain history. But you can also plan an MPS using the planning process in either Master Scheduling/MRP or Advanced Supply Chain Planning to generate an MPS. This process applies the netting logic described earlier to calculate net requirements and plan orders for those items you have designated as Master Scheduled.


**NOTE**

It is possible to load an MPS with forecasts and sales orders; in this case, the MPS would function exactly like an MDS. However, this usage obscures the purpose and distinction of the two schedule types; we will assume throughout this book that you will use an MDS to represent independent demand and an MPS to represent a production plan.

Master Schedule Load

The process of copying data into a master schedule is called a *load*. Because the master schedule represents the demand that drives the planning process, you must reload the master schedule as often as you need to recognize changes in the demand. The process is simple, but it does provide a few options:

- Name of the schedule you are about to load.
- Source of the information (source type, source organization, and source name).
- If you want to include Sales Orders in the master schedule, you can choose to load All Sales Orders or Sales Orders from the Start Date Forward. Additionally, you can limit sales order lines to a specific demand class only.
- The treatment of the Demand Time Fence. You can choose to Load Forecast outside the demand time fence only. If you are including sales orders and have consumed the forecast, this is a typical choice; it results in sales orders only inside the demand time fence and the greater of sales orders or forecast outside the demand time fence. (The demand time fence is defined as an attribute of each item.) You can also choose to Load orders within and forecast outside demand time fence (ignoring any large orders outside the demand time fence) or to Ignore the demand time fence (loading all specified data regardless of the demand time fence).
- A start and end date, to limit the information you will load.
- You can choose to explode the forecast when you load a master schedule.
- You can load current (consumed) forecast quantities or original (unconsumed) forecasts, and you can consume the forecast if you have not already done so. These options are the key to utilizing two-level master scheduling, which is discussed later in this chapter.
- The overwrite option: All Entries, None, or Same Source Only.

**NOTE**

The Overwrite option first deletes any information that matches the overwrite criterion; the load process then re-creates the appropriate data. As with forecasts, you must choose the Overwrite option carefully, or you may duplicate existing data. Same source refers to data with the same type and name. For example, if you were loading a specific forecast into a master schedule, the Same Source Only option would delete any existing forecast data with the same type (forecast) and the same forecast name; any data from different forecasts would not be affected.

Several other options are useful primarily for simulations:

- **Modification** Modifies the loaded quantities by a positive or negative percentage.
- **Carry Forward Days** Shifts the quantities forward or backward in time. Enter a positive number to shift the quantities forward; a negative number to shift backward.

Source List

If you have many forecasts, master schedules, or plans that you will regularly consolidate into one master schedule, you can define a source list that identifies all these objects with a single name. Then, when loading a master schedule, you can specify the single source list name, rather than running multiple loads for each forecast, schedule, or plan you want to load.

Open Master Schedule Interface

You can populate master schedules using the Open Master Schedule Interface. Like the Open Forecast Interface, described earlier, you must first define your items, organizations, and master schedule names. Then, load your data into the table MRP_SCHEDULE_INTERFACE. Required data includes

- INVENTORY_ITEM_ID
- SCHEDULE_DESIGNATOR
- ORGANIZATION_ID
- SCHEDULE_DATE

- SCHEDULE_QUANTITY
- PROCESS_STATUS

Set the PROCESS_STATUS to 2 to designate data waiting to be processed.

This process is also controlled by the Planning Manager; it periodically checks the MRP_SCHEDULE_INTERFACE table and launches the Master Schedule Interface Load Program to process new data it detects. The load program validates the data in the interface table and loads valid data into the specified master schedule. Errors are handled just like the Open Forecast Interface; the load program sets the PROCESS_STATUS to 4 and enters a text message in the ERROR_MESSAGE column in the interface table. View and correct invalid data with SQL*Plus or a custom program.

Reviewing the Master Schedule

Oracle provides several forms and reports where you can review your master schedules. Online, you can use the Item Master Demand Schedule Entries and the Item Master Production Schedule Entries forms to view master schedule data; you can view the schedule details in the same form you entered them or consolidate the data into weekly or period buckets. The Master Schedule Detail Report provides detailed information, in a bucketed horizontal display, a detailed vertical display, or both. The Master Schedule Comparison Report lets you compare two master schedules.



NOTE

The Master Schedule Detail Report, and many other bucketed reports in Oracle's planning applications, let you control the level of detail reported. The reports are formatted to show multiples of 12 buckets; you choose the number of buckets (12, 24, or 36) that you want to see on the report. Then you specify how many of those periods are weeks; the remaining buckets will show "period" information. To ensure that the period buckets show data for a complete period, the number of weekly buckets you can choose is determined by your planning calendar and the number of weeks remaining until each period start date. For example, if you use a 4-4-5 calendar for planning buckets, and it is week 1, you could choose 4, 8, or 13 weeks; if it is week 2, you could choose 3, 7, or 12 weekly buckets.

Schedule Relief

Just as forecasts must be consumed to avoid overstating demand, master schedules must be reduced to prevent overstating demand or supply. If a sales order MDS entry remained after the sales order had been shipped, you would overstate demand. Similarly, if an MPS entry remained after you had created a discrete job or purchase requisition, you would duplicate supply. The process of removing this duplication is called *relief*, and it is another function of the Planning Manager.

Shipment Relief

Because an MDS represents a statement of expected shipments, it is relieved when the demand no longer exists (i.e., when a sales order has been shipped). As long as the Planning Manager is running, the process is automatic.

Production Relief

Master Production Schedules are relieved when a real statement of production replaces the plan that the MPS represents. Note that this is not when the product is actually produced; it is when the production order (a WIP discrete job or repetitive schedule) is created. While this strikes some people as unusual, it is quite sensible— if you had both an MPS entry and a WIP job, you would be doubling your statement of production. Again, the process is handled by the Planning Manager and is automatic.

Two-Level Master Scheduling

If you use Product Families to forecast, you have the option of utilizing two-level master scheduling. This technique lets you determine how remaining aggregate forecasts will be exploded and used in the planning process and further refines your control over the forecasting and planning process.

To understand two-level master scheduling, consider the following simple example in Figure 8-6. The bill of material structure illustrated shows an aggregate item, A, with components B and C, distributed at 50 percent each. A forecast for 100 A, therefore, generates a forecast for 50 each of B and C.

If A is a Planning Bill item, two-level scheduling is not supported, and the forecast is consumed only at the components, B and C. A Sales Order for 60 of item C will consume the entire forecast for C, actually overconsuming C's forecast. But the demand for B remains unchanged. Thus, if you were to run a plan with all the current demand data, the total demand is 110 units, even though the aggregate forecast was only 100.

In the two-level scheduling scenario, the forecast is consumed at both the aggregate and the component levels. You can choose to re-explode the remaining aggregate forecast, with the results shown in Figure 8-7. Note that this will keep the aggregate demand constant, until the entire aggregate forecast is oversold. Notice, too, that this dynamically adjusts the mix of B and C, based on the forecast consumption.

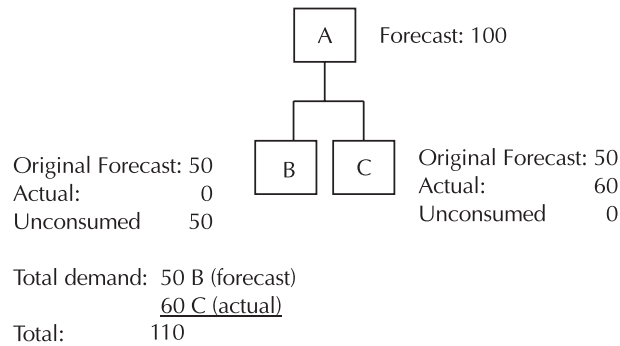


FIGURE 8-6. *Traditional forecast consumption and explosion*

The choice of forecast explosion methods is controlled by the explosion options used when you load the master schedule from a forecast. The Load/Copy/Merge Master Demand Schedule form lets you copy forecasts into your master schedule. Options let you choose whether to load the current (i.e. the remaining forecast, after consumption) or original forecast quantities, whether to explode the forecast, and whether to consume the forecast if consumption has not already occurred. If you explode the original quantities and select the Consume option, you will get the demand shown in Figure 8-6. If you explode the current quantities, you will get the demand picture that we've been calling two-level master scheduling (see Figure 8-7).

There are several potential pitfalls to be aware of in this process. If you explode original quantities, include sales orders, and do not select the consumption, you will inflate demand. The load process will do exactly what you ask—it will explode the

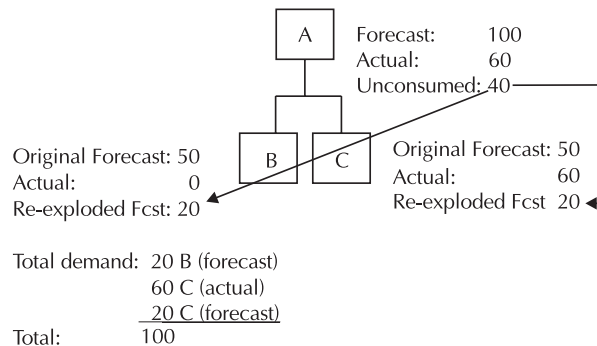


FIGURE 8-7. *Two-level master scheduling*

original quantities without consumption and add in sales orders that otherwise should have consumed that forecast. This will result in the very duplication of demand that forecast consumption is designed to prevent. (The load process does not allow you to consume if you are loading the current, i.e., already consumed, quantities.)

A second pitfall is loading a forecast that you have already exploded. This also results in duplicate demand—you will load the existing forecast for the components or family members and then create additional forecast by exploding the family forecast.

Capacity Planning

Oracle provides two types of capacity planning within its planning products: Rough Cut Capacity Planning (RCCP) for evaluating the high-level capacity requirements of a master schedule and Capacity Requirements Planning (CRP) for evaluating the detailed capacity requirements of an MRP (or ASCP) plan. Both types of capacity planning compare the available capacity with the calculated requirements. Available capacity comes from your resource definitions, as described in Chapter 4; what differentiates CRP from RCCP is the method of calculating the resource requirements.

Capacity Requirements Planning (CRP)

CRP calculates required capacity from the routings used to build your products. It is always run as part of running an MRP (or DRP) plan, since it needs to know the planned orders for its calculation.

In calculating resource requirements, CRP uses the primary routing for planned orders, but will use alternate routings if they have been specified for existing jobs. It recognizes modifications you have made to WIP resources (for example, specifying the use of an alternate resource). Because it is run as part of MRP, it benefits from the netting logic inherent in planning—if your on-hand quantities don't require planning of new replenishment orders, you won't generate capacity requirements. And if jobs on the shop floor are partially completed, CRP recognizes which resources have already been used and only plans requirements for the resources that remain.

Rough Cut Capacity Planning (RCCP)

RCCP lets you check your master schedule against critical resources that you identify in a bill of resources. RCCP can be used to evaluate either master production or master demand schedules. In a manufacturing environment, it is more typical to use RCCP to evaluate production schedules because the constraints are most often manufacturing resources. In a distribution environment (or for a quick check on distribution-related resources), you can use RCCP to evaluate master demand schedules.

What makes rough cut rough? Since it does not use detailed routings, it is inherently less detailed. Though it might seem obvious, a hidden implication is in the requirements for subassemblies. Because RCCP does not depend on the netting and explosion logic of MRP, it may show resource requirements for subassemblies

even if a subsequent MRP plan will find that there are no net requirements because of on-hand quantities. And there is no visibility to the progress of WIP jobs.

Bills of Resource RCCP calculates resource requirements simply by multiplying the master schedule quantity by the resource requirements expressed in a bill of resources that you define for the master scheduled item. In a manufacturing environment, you can construct a bill of resources by rolling up the routings; this process is called a bill of resource load. The result might be considered a kind of “aggregate routing” for the item; it lists all the resources you need to manufacture the master scheduled item and all its subassemblies. You can then modify the bill of resources by deleting noncritical resources or by adding resources that are not normally part of your routings—shipping capacity, for example.

Strictly speaking, a *bill of resources* is a listing of the resources you want to evaluate for a single item. But, to be useful, any evaluation of capacity must look at all the demands placed on your resources. It doesn’t do much good to know that item A requires 20 hours of a resource, unless you also know the demands placed by items B and C—it’s the total resource requirements that you must compare to the resource availability. For this reason, bills of resource are grouped into a bill of resource set; when you evaluate rough cut capacity planning, you must specify which bill of resource set you want to use to calculate resource requirements. Like so many objects in the Oracle Applications database you can have multiple bill of resource sets; this might be useful for simulations, or to evaluate distribution resources against your demand schedules and manufacturing resources against your production schedules.

Once you have defined a bill of resource set, you can view your rough-cut capacity plans on demand; the calculations are done when you perform the inquiry or run the report.

Use the Load Bill of Resource concurrent program to calculate bills of resource from your routings. Parameters include

- The bill of resource set you are loading.
- Whether you want to preserve any manual additions or modifications you may have made to the bill of resources, if it already exists.
- A resource group, to limit the creation of the bills of resource to items within a specific resource group. (Resource groups are discussed in Chapter 4.)
- Whether you want to create a bill of resource for all selected items or for Model and Option Class items only.
- How you want to deal with other master scheduled items that might be structured into bills of material at lower levels.

This last option, labeled Rollup MPS Items on the Parameters window, deserves explanation. Consider the example in Figure 8-8; if the only master scheduled item were A, you would typically want A’s bill of resource to show two hours of labor—one hour to assemble C and one more hour to put B and C together to make A. But if C were also a master scheduled item, you would not want to roll its time up into A’s bill of resource because you could be overstating the resource requirements. If your master schedule includes both A and C, a rough cut capacity plan would calculate one hour per unit for C and two more hours for A. This would overstate the capacity requirements—one of the two required hours was already planned for by C’s bill of resource, and only one more hour of labor is needed to assemble product A.

You avoid this problem by setting the parameter Rollup MPS Items to No; this “breaks” the rollup process at each master scheduled item and ensures that a bill of resources does not contain resources that were already accounted for on other master scheduled items lower in the bill of material. Note that the parameter does default to Yes. While it’s possible to construct a scenario where this could make sense (e.g., you plan to do rough cut capacity planning only against your top-level items), it would be inappropriate if you plan to do rough cut planning for all of your master scheduled items. Consider this carefully when loading bills of resources.

You can also build bills of resource “from scratch” by listing the resources you want to check in your RCCP plan. If you define bills of resource manually in a

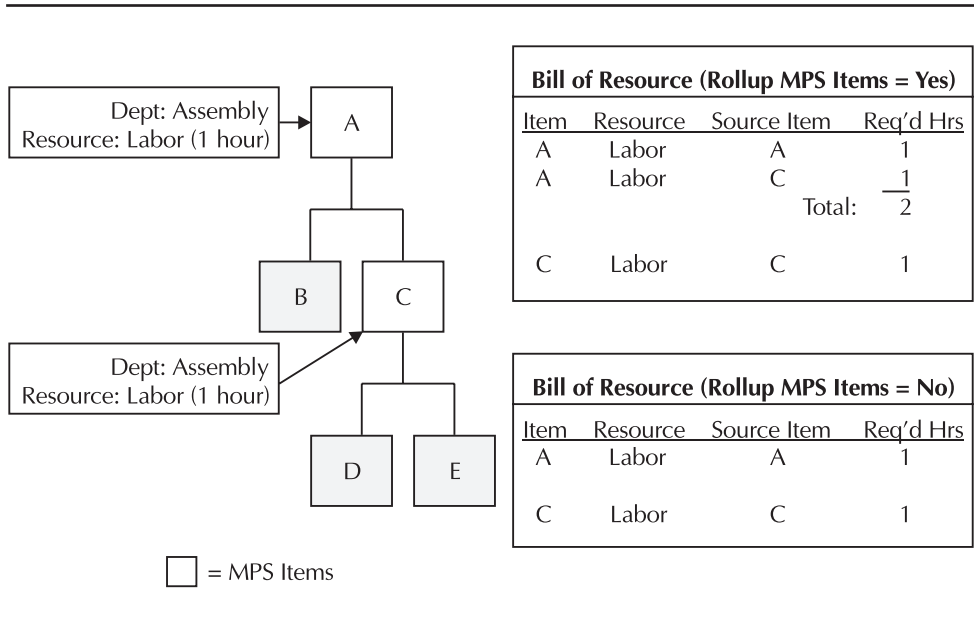


FIGURE 8-8. MPS rollup options

manufacturing environment, you may want to include the critical resources you need to build subassemblies of the master scheduled item.

A bill of resource for an item consists of one or more rows of data that specifies the following attributes for each resource you list. On the Main tab, enter the following information:

- **Department or Production Line and Resource critical to the production or distribution of the item** These are often the same departments and resources you use in your production routings, but you can also define other resources that might not traditionally be part of a routing. For example, if shipping capacity is critical to your business planning, you could define it as a resource within a department and use RCCP for a quick evaluation. The key thing to remember is that you must define these resources and assign them to departments, defining their available capacity just as if you were going to use them in a routing. This process is described in Chapter 4.
- **Source Item** That is, the item responsible for the demand. In a rolled-up bill of resources, this will be the item whose routing contributed the resource usage, either a subassembly or the item itself.
- **Setback Days** The number of days prior to the schedule date that the resource is required. For example, if the resource is required two days prior to the schedule date of the item, enter **2**.
- **Usage Quantity** The exploded quantity from the bill of material. This is the quantity of the source item required to make the item whose bill of resource you are defining; if the source item is the item itself, this quantity will be 1. Note that this is not the quantity of the resource required; you enter that information on the Discrete Resource tab.

On the Discrete Resource tab, enter the following information:

- **Required Hours and Basis Type** Just as on a routing, a Lot basis means that the resource hours are required only once per lot (in this case, master schedule entry); for example, a setup. An Item basis means that the required hours are need for each unit of the item on your master schedule.
- **Op Seq** The operation sequence where the resource is used. This will be populated for you by the bill of resource rollup; if you are defining (or modifying) a bill of resource manually, you may specify the correct operation or enter **1**.

The bottom of the form displays the source of the entry: Rollup means that the entry is the result of a routing rollup. Manual addition or Manual update indicates

that you have created or modified this entry manually; this distinction allows you to re-roll routings and optionally preserve manual modifications to your bills of resources.

Planning for Engineering Changes

A key element of planning is the ability to recognize anticipated bill of material and routing changes. This allows you to phase in changes to the products you manufacture and have planning react appropriately—planning will stop generating orders for the obsolete components or resources and begin ordering the new components or resources.

Effectivity Dates

Both bills of materials and routings use effectivity dates, indicating when a component or routing operation begins to be effective, and optionally when the component or operation is disabled. Effectivity is determined by the start date of a planned order; planning (and WIP, if you manually create a job or schedule) will select only those components that are effective on the date the job is scheduled to start.

The effective date, labeled “From,” defaults to the current date when you create a bill or routing; you can select a future date if the element is not yet effective. You cannot select a date in the past. If the ending date, labeled “To,” is left blank, it indicates that there is no planned end to the effectivity of the component or operation.

Note that the To date is a disable date, the date (and time) that the component or operation is no longer used. This is important if you are replacing one component with another; you set the starting (From) date of the new component equal to the disable (To) date of the old component so that the new component becomes effective at the same time that the old component is discontinued. Unlike some systems that might term the end date an Effective Thru date, if you were to disable a component on the 15th, for example, but not start the replacement until the 16th, a job or planned order that started on the 15th would contain neither component.

When you’re setting the effective dates for a new component, it’s helpful to know its Total cumulative lead time. If you plan to add a component within its cumulative lead time, you might have to expedite orders in order to obtain the component in time.

Effective dates can be entered manually on bills or routings; effective dates on bills of material can also be maintained with Engineering Change Orders (ECOs). With Release 11i.5, you can also use ECOs to modify routings.

Engineering Change Orders (ECOs)

ECOs provide a more controlled environment for maintaining effectivity dates than the simple bill of material maintenance process. ECOs (described in Chapter 5) let

you group multiple changes under a single control number. With an ECO, you set an effective date for all of the changes you want to make to an individual bill of material; this eliminates any possibility of “gaps” if you define a replacement component and lets you keep multiple component changes in synch with one another. ECOs allow the system to calculate the effectivity dates for a set of changes based on using up a component on a bill (or using up the assembly itself). ECOs let you control whether a set of changes is visible to the planning applications. And ECOs can base the schedule date of a change on the date on which an assembly or component is scheduled to be used up.

Schedule Date Changes

As you change effectivity dates, planning will replan accordingly. Moving a schedule date in (scheduling the change to occur earlier than originally planned) can result in expedite suggestions from planning; moving a date out could result in shortages of the old component, particularly if you have deactivated it.

If a pending change is driven by the use-up of the assembly itself, the schedule date of the change will be updated automatically if MRP calculates a new use-up date for the assembly. If the change is driven by the use-up of a component of the assembly, MRP will trigger the ECO Use Up Alert to notify the appropriate planner.



NOTE

At present, Use-Up planning is only integrated with MRP, not with Advanced Supply Chain planning.

Discrete versus Repetitive Planning Logic

Oracle Applications allow you to plan for both discrete and repetitive production methods in a single plan. Discrete production involves recording production against discrete jobs, sometimes called *work orders* or *production orders*. Discrete jobs represent production of a specific quantity of an item, completed on a specific date. Repetitive production records production against repetitive schedules, which represent a continuous rate production for a period of time. While much of the planning logic is the same, there are a few key differences between discrete and repetitive production that necessitate differences in the planning logic.

NOTE

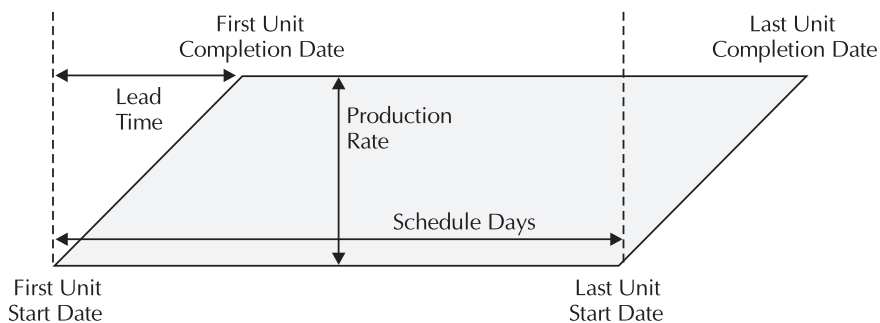
Don't confuse repetitive production with flow manufacturing. Repetitive production is designed for relatively high volume, continuous production of standard items. Flow manufacturing is designed for rapid production of individual items, even custom-configured products, at varying rates of production.

Repetitive Production Concepts

A repetitive schedule represents continuous production over a period of time; it models an ongoing rate of production and ongoing consumption of components. Compare this with a discrete job that represents production of a specific quantity of a product on a specific date and models an individual requirement date for each component. Whereas a discrete job has a start date and a completion date, a repetitive schedule is defined by four dates:

- **First unit start** The date that the first unit is started on a production line
- **First unit completion** The date that the first unit is completed on the end of the production line
- **Last unit start** The date the last unit on the schedule starts down the production line
- **Last unit completion** The date that the last unit is completed

These dates are represented in the following diagram:



Discrete jobs have a total job quantity to be completed on a specific date. Each component on a discrete job is required on a specific date. Repetitive schedules, by contrast, represent a rate of production for a period of time (defined by the first and last unit completion dates). On a repetitive schedule, products are produced and components are required continuously over the duration of the schedule.

The nature of repetitive production requires slightly different planning logic than that described earlier; the following pages discuss differences between planning for discrete production and repetitive schedules.

Repetitive Planning Logic

The objective of repetitive planning is to calculate a smooth rate of production for a period of time. When planning repetitive production, the process calculates the rate of production by taking the average daily demand within user-defined buckets. This process requires some unique planning logic and uses some terminology that has not been discussed previously.

Planning Periods

You establish planning periods on the Planning Parameters form for each inventory organization. You define an Anchor Date, which marks the start of your repetitive planning periods. Initially, this should be the start date of the first repetitive planning period; after that, each time you run a plan the system will roll the anchor date forward, if necessary, to keep the anchor date consistent with the start of the first repetitive planning period. You also define two horizons as a number of days; this breaks up the planning horizon into three sets of periods. Within each, you define the size of the buckets. Smaller buckets are often used in the first set of periods, so that repetitive planning is more reactive to changes in demands, while larger buckets in the later sets of periods allow planning to generate smoother production plans.

When you define repetitive planning periods, you can choose whether to use workdays or calendar days. Though it seems counterintuitive, it is generally advisable to use calendar days to determine the length of your planning periods. Using calendar days lets you ensure that your planning periods always start on the same day of the week, regardless of holidays or other non-work days. Consider Figure 8-9; if you define your repetitive periods to start on a Monday and define each period as five work days in length, the first scheduled non-work day will shift the next period to start on Tuesday.

If instead you define your repetitive periods as seven calendar days long, each period will start on the same day of the week regardless of intervening holidays, as shown in Figure 8-10. Calculation of average daily demand will always use the number of work days within a period, even if the length of that period was determined by using calendar days.

Day	1	2	3	4	5
Demand			40		10
Production Rate	10	10	10	10	10

FIGURE 8-11. *Average daily demand*

An implication of this planning method is that demand is generated daily for the components of a repetitive schedule. This generally results in a large number of discrete demands. If you are not prepared to respond to these individually, you should consider adding order modifiers, for example, a minimum order quantity or fixed days supply.

Recognition of Existing Schedules

Perhaps the most unconventional aspect of repetitive planning is that it does not recognize, nor suggest changes to, any existing schedules. It always suggests new schedules. In part, this is because repetitive schedules have no identifying number (like a discrete job); they are identified by the rate and the four key dates. Suggesting a change to any of these parameters, therefore, means suggesting a new schedule. Thus, repetitive planning does not need to know the existing schedules; its job is to suggest the right rate of production.

Another implication of the fact that repetitive planning does not consider existing schedules is that any changes you made to the component requirements of a repetitive schedule are not considered in planning. If you have made a component substitution, for example, the planning process does not recognize it. Thus, you must carefully monitor the availability of substitute material (or excess plans for the original material). While this may appear problematic, remember that repetitive production is designed for high-volume production of standard items; if you are making frequent changes to component requirements, the product may not be a good candidate for repetitive manufacturing.

Treatment of Order Modifiers

Like discrete planning, repetitive planning uses order modifiers—minimum and maximum order quantities and fixed order quantity. While these are the same modifiers that discrete planning uses, they are interpreted as modifiers to the *daily rate*, not to the total schedule quantity. Keep this in mind if you use modifiers on repetitively planned items, or you can dramatically inflate your schedule rates.

**NOTE**

One exception to the use of order modifiers is the Fixed Days Supply modifier. Because repetitive planning inherently plans for a fixed number of days (as determined by your repetitive planning buckets), the Fixed Days Supply modifier is simply ignored for repetitive items.

Planning Controls

To allow you to build stability into repetitive plans, Oracle provides several controls that determine the maximum increase or decrease in repetitive schedules at different points in time. Oracle also allows you to specify how much overproduction you will tolerate, in order to reduce suggested changes to production rates.

First, within an item's lead time, planning will not suggest any new repetitive schedules. This reflects the fact that it is often difficult or impossible to change the production rate once an item has started down the assembly line.

Second, within an item's planning time fence, Oracle makes use of the item attributes Acceptable rate increase percent and Acceptable rate decrease percent. These attributes let you specify how much of a change is allowed within the planning time fence. These percentages are applied to the current rate of production and determine the maximum and minimum rates that planning will suggest for new schedules. For example, if the current production rate is 100 units per day, an Acceptable rate increase of 10 percent means that planning cannot suggest a new rate greater than 110 per day within the planning time fence. Similarly, an Acceptable rate decrease of 10 percent means that planning could not suggest a rate lower than 90 per day. Outside of the planning time fence, however, planning is free to suggest whatever rate is required.

Note that zeros and nulls are treated very differently in these attributes. A zero means the acceptable increase or decrease is indeed zero; no rate changes are allowed. Leaving the attribute blank (null) implies that there is no limit on the acceptable increase or decrease; planning can make whatever suggestions it wants.

To dampen minor fluctuations in repetitive schedules, planning uses another item attribute, Acceptable overrun percent. This is the amount that the current daily rate may exceed the optimal rate before planning will suggest a new rate. In other words, it is the rate of overproduction that is acceptable in order to avoid minor changes in scheduled rates. The overrun percent is applied to the optimal production rate, and the result is compared to the current aggregate rate. For example, if the optimal production rate is 100 units per day, and the Acceptable overrun percent is 10 percent, your current production rate could be as high as 110 per day before planning would suggest a new schedule with a lower rate.

Repetitive Schedule Allocation

Repetitive planning begins with the calculation of the total production required; this may be accomplished on multiple production lines, so the total required is referred to as the aggregate repetitive schedule. Once the aggregate rate of production is known, planning will allocate production across all the lines designated for the product. When you define production lines for an item, you specify the maximum rate the line can produce and assign a priority to that line and item combination; planning uses these rates and priorities to allocate production across lines.

The allocation process first attempts to assign production to the highest priority line (lowest number) for the item. Allocation will load the first line to its stated capacity, then load the second line to its capacity, and so on. If there is not enough capacity across all the designated lines for an item, allocation will load the excess production onto the first line. Note that while allocation will attempt to respect the maximum stated capacity of a line for an item, it does not consider the load placed by other items that may also be scheduled on the line. Also keep in mind that it will overload the primary line for an item, if necessary.

Repetitive Schedule Implementation

You can implement repetitive schedules from the Planner Workbench in much the same way that you implement discrete jobs; simply select the job or schedule for release, and initiate the load process from the Planner Workbench. This initiates the WIP Job/Schedule Interface program to load the job or schedule into WIP. But while a discrete job may be loaded in either a released or unreleased status, repetitive schedules are always loaded with a status of Pending – Mass Loaded.

This status avoids a restriction in WIP on the existence of overlapping repetitive schedules. WIP will not allow schedules for the same item on the same line to overlap, unless the schedules have a status of Cancelled or Pending – Mass Loaded. As mentioned earlier, planning always suggests the needed schedule and does not suggest changes to existing schedules. If you have already implemented a repetitive schedule for the same item, line, and time period, a new schedule would overlap. Thus, schedules implemented from the Planner Workbench always carry the status of Pending – Mass Loaded. You must release these schedules in WIP, respecting the restriction on overlapping schedules, before you can report production against the new schedule.

When you release a repetitive schedule, you have the opportunity of closing or canceling existing schedules; this lets you avoid overlapping schedules. Alternatively, you can simply modify the rate or dates of the existing schedule based on the suggestions of the new schedule; then you can cancel the new schedule. This is preferable if you have issued material or charged labor to the existing schedule, as it avoids misleading variances. The use of the Planner Workbench is discussed in

detail in Chapter 11 and 12; repetitive production is described in more depth in Chapter 16.

Configuring Repetitive Planning

Repetitive planning requires the same setup as discrete planning. In addition, because of its unique requirements, it requires the following:

- Define production lines on the Production Lines form.
- Associate items and production lines using the Repetitive Line/Assembly Associations form.
- Define repetitive planning buckets on the Planning Parameters form.
- Designate the desired items to use repetitive planning logic on the Master Items or Organization Items forms.
- Set repetitive planning controls, also on the Master Items or Organization Items forms.

Kanban Planning

Another type of planning that Oracle Applications provides is kanban planning, part of Oracle's support for Flow Manufacturing techniques. As noted in Chapter 1, *kanban* is a Japanese word that suggests "sign" or "signal." From a manufacturing perspective, a kanban is a visual indicator—sometimes a card, sometimes an empty bin or rack—that triggers replenishment of an item. The intent is to streamline the manufacturing and procurement process; rather than relying on a computer simulation to suggest when to order, kanban replenishment triggers orders based on an operator actually signaling that more material is needed.

While kanban replenishment does not rely on MRP planning to launch the replenishment orders, effective use of kanbans does require appropriate sizing. If a kanban is too small, you might run out of material before it is replenished; if it is too big, you carry excess inventory. Oracle provides a kanban planner to size kanbans, and to resize them when necessary. In addition, though kanbans eliminate the launching of orders from MRP, MRP planning can still be valuable—it can provide calculation of usage over time, which can be useful in supplier negotiations or when deciding to expand manufacturing capacity.

Kanban Setup

You designate items as kanban-controlled through the item attribute Release Time Fence. Setting this attribute to Kanban Item (Do Not Release) prevents you from releasing orders from the Planner Workbench.

For the purposes of planning, kanban setup involves defining pull sequences, which specify the source of replenishment and planning parameters for kanban-planned items. For each item under kanban control, you define the subinventory and optional locator that you will replenish with the kanban. Then define the source of the replenishment (intra-org, inter-org, supplier, or production) and the appropriate source information. For example, for an inter-org kanban, you would specify the source organization, subinventory, and optional locator from which the material will be pulled.

The factors that control planning are the planning method, lead time, and lot sizing information for the specific kanban location. There are three planning methods:

- **Kanban size** With this method, you specify the number of cards or kanbans you desire, and the planning process calculates the size of those kanbans, i.e., how much material must be replenished each time the kanban is signaled. For example, if you elect to employ a simple “two-bin” replenishment system, you would specify 2 as the number of kanbans; planning would calculate how much material is required in each bin to keep the line flowing smoothly and efficiently.
- **Number of cards** Using this method, you specify the size of the kanban, and planning calculates the number of kanbans you need. For example, your kanban might be a rack or tote capable of holding a fixed amount of material, or you might determine the container size that an operator could move safely without assistance; planning determines how many of those containers you need.
- **Do not calculate** With this method, you size your kanbans manually. This method excludes an item from kanban planning, while still letting you use the kanban to signal replenishment.

In addition, there are lead times and lot sizes defined for each pull sequence that affect the planning process:

- **Minimum order quantity** This is defaulted from the item, but you can override for a specific pull sequence.
- **Lead time** This is the lead time to replenish one kanban from its source.
- **Lot multiplier** Use this if it makes sense to replenish in fixed multiples—for example, if the item has a natural lot size.
- **Allocation percent** This represents the distribution of an item’s demand across multiple locations.
- **Safety stock days** The number of days of safety stock you want to maintain.

Additional detail on kanban setup is provided in Chapter 15.

Kanban Calculations

Kanban planning calculates the required number or size of kanbans, based on the planning method you specified for each pull sequence. Basically, the process looks at the average daily demand you supply and the planning parameters from each pull sequence and determines how much material or how many kanbans are needed to supply that demand. The details of the calculation program are described in detail in the Oracle Master Scheduling/MRP and Oracle Supply Chain Planning User's Guide, part of the Oracle documentation library provided with the software.

Average Daily Demand

A key to kanban planning is the calculation of average daily demand. When you define a kanban plan, you specify the source of demand—you can use a forecast, a master schedule (MDS or MPS), or actual production. For your initial kanban plan, you might want to use the same forecast or master schedule used for validating your line design on the Mixed Model Map (see Chapter 7). This forecast or schedule should be typical of the mix and volume of demand for which you are designing your line. On an ongoing basis you should check your kanban sizes against actual or planned production; this will give you early warning signs of changes in demand that might require intervention—perhaps the creation of a non-replenishable kanban to accommodate a one-time spike or a shift in demand patterns that might require a permanent change in your kanban sizes. This ongoing evaluation is often overlooked in flow manufacturing environments, but the simplicity of the Kanban Planner makes it easy to do.

Once the demand is identified, the kanban planning process performs an explosion much like MRP to determine dependent demands for your kanban items. The planner then applies the factors you defined in your pull sequences and uses the calculated yields from your flow routings to determine the optimal size or number of kanbans.

Effectivity Date Changes

It is important to note that kanban planning does not take into account effectivity date changes. The explosion process determines average daily demand using the effectivity date you specify when you run the plan; it does not see pending bill of material changes. Because kanban planning uses the average daily demand, recognizing pending bill of material changes would result in incorrect kanban sizes. For example, if you were replacing one component with another, the average demand would be too low for each component, and the kanbans would be sized inappropriately.

Generating and Implementing Kanban Plans

To generate a kanban plan, you first define one or more plan names using the Kanban Names window. The only parameter you need to specify is the source of the demand you want to use.


Then, launch the plan on the Launch Kanban form. Specify the effectivity date to use when exploding bills of material. You can optionally limit the plan to a range of items or item categories, and you can enter start and cutoff dates to restrict the demand data the plan considers.

When the planning process is complete, you can view the results on the Kanban Workbench, shown in Figure 8-12. The Details button shows the results of the planning process and the kanban parameters upon which that plan is based. You can view the demand used by the Kanban Planner with the Demand button.

If you want, you can rerun a portion of the kanban plan with new parameters—select one or more kanbans by checking the box to the left of each row on the Details window, and modify the desired parameters. You can modify the type of result you want to calculate (kanban size or number of cards), as well as source information, minimum order quantity, lead time, fixed lot multiplier, and safety stock days. For example, if you found that some kanbans were sized with odd quantities, you could specify the appropriate fixed lot multiplier. Save your changes, and click the Simulation Replan button. This will submit a request to run the Kanban Planner; when the planner completes, you can query the form to see the new results.


Item	Subinventory	Locator	Variance (%)	Plan 1 Cards	Plan 1 Size	Plan 2 Cards	Plan 2 Size
<input type="checkbox"/> KB15138	RIP		66.67	25	1	25	3
<input type="checkbox"/> KB15138	Stores	1.1.1..	53.57	13	10	28	10
<input type="checkbox"/> KB18759	RIP		33.33	2	50	3	50
<input type="checkbox"/> KB18759	Stores	1.1.2..	33.33	100	2	100	3
<input type="checkbox"/> KB28287	RIP		0	2	100	2	100
<input type="checkbox"/> KB28287	Stores	1.1.3..	42.86	4	100	7	100
<input type="checkbox"/> KB42047	RIP		33.33	2	50	3	50
<input type="checkbox"/> KB42047	Stores	1.1.4..	0	2	1000	2	1000
<input type="checkbox"/> KB86324	RIP		0	5	500	5	500

FIGURE 8-12. View planning results on the Kanban Planner Workbench

**NOTE**

Though Oracle calls this a simulation, there is no way to simply reset the plan to its state prior to the replan. To reset, you can either run the Kanban Planner again or change the parameters back to their original values and run the “simulation” again.

You can compare different plans, or you can compare the plan with the values that are currently driving your production execution system. When you press the Compare button, the Find window will ask for the name of the comparison plan; by default, it will select Production (i.e., the current definitions from your pull sequences), but you can select an alternative plan if you want.

**TIP**

To filter out trivial or nonexistent changes between plans, consider including selection on the Variance field in the Find window. For example, including the criterion Variance Greater Than 0 will show only differences between the plans.

You can implement suggested changes from either the Kanban Details or the Compare Kanban Plans window. To implement a change, select one or more kanbans with their selection check boxes (or use the Select All button), and click the Update Production button. This will update your current pull sequences with the information from the selected kanban plan. You will need to regenerate and reprint the kanban cards as described in Chapter 15; for convenience, the Tools menu on the Kanban Planner Workbench provides a link to the Kanban Cards form and allows you to print selected kanbans.

Planning Parameters and Profiles

A number of parameters, profiles, and item and bill of material attributes control the planning process. These are described in detail in the user guides for Master Scheduling/MRP, Inventory, and Bills of Material, part of Oracle’s documentation library. A few highlights are noted here.

Planning Parameters

The Planning Parameters form lets you specify execution defaults and repetitive planning parameters for your MPS, MRP, and DRP plans in each inventory

organization. These parameters are copied to each plan you define, but you can override them in the Plan Options window.

MRP Profiles

There are almost 100 profile options that control the execution and defaults of MRP plans (number of workers, debug mode, etc.). Most of these can be set at the site level, but several may be useful to set defaults for individual users. For example, the profiles MRP: Default Forecast Name, MRP: Default Plan Name, MRP: Default Schedule Name, and MRP: Default DRP Plan Name provide defaults for the forecasts, MRP plans, master schedules, and DRP plans that a user will work with on various inquiry and update forms. A default plan name will open automatically when you invoke the appropriate form; this saves you from having to find or query the plan you want to use (though you can override the default if desired).

Item and BOM Attributes

Much of the detail of planning is controlled by the various item and bill of material attributes. These attributes were mentioned briefly in earlier chapters, but several key attribute groups (corresponding to tabs on the Master Items and Organization Items forms) are especially important to the planning process. These are described in the following sections.

General Planning Attributes

The General Planning tab (shown in Figure 8-13) includes attributes that control MRP and ASCP planning, as well as the simpler planning methods provided in Oracle Inventory (Min–Max and Reorder Point Planning). Of special interest are the following:

- **Planner** This field identifies the planner responsible for the item. Besides being a mechanism to select items in the Planner Workbench and sort various planning reports, the planner is required for ASCP to generate recommendations. The planner also must approve any move orders requesting the item. Define planner codes with the Planners form; optionally specify their employee name and e-mail address.
- **Make or Buy** This field, along with sourcing rules (described in Chapter 11), controls the default order type planned for the item. If it is set to “Make,” planning explodes the item’s bill of material and suggests discrete jobs to satisfy demand; if it is set to Buy, planning will not explode a bill of material and will suggest purchase requisitions. You can override the suggested order type based on the setting of the Purchasable and Build in WIP

The screenshot shows the 'Organization Item (V1)' window with the following details:


- Organization: V1 Vision Operations
- Item: MX344523
- Description: Mother Board
- Display Attributes: Master (unselected), Org (selected), All (unselected)
- Physical Attributes | **General Planning** | MPS/MRP Planning | Lead Times | Work In Process | Order Management
- Inventory Planning Method: Not Planned
- Planner: B. Daniels
- Make or Buy: Make
- Min-Max Quantity: Minimum [], Maximum []
- Order Quantity: Minimum [], Maximum []
- Cost: Order [], Carrying [] %
- Source: Type [], Organization [], Subinventory []
- Safety Stock: Method: Non-MRP Planned, Bucket Days [], Percent []
- Fixed Order Quantity [], Fixed Days Supply [], Fixed Lot Multiplier []

FIGURE 8-13. *Item general planning attributes*

attributes. (This field also determines the default of the Based on Rollup attribute in Oracle Cost Management.)

- **Order Quantity Modifiers—Minimum, Maximum, Fixed Order Quantity, and Fixed Lot Multiplier** These attributes determine the quantity of the planned orders that MRP or ASCP can suggest; use them if there is a natural or economic lot size for the item (e.g. certain products are not practical to order in small or odd quantities; others might require a separate order if the order would be more than a full truckload, etc.). But use these with caution in a “lean” manufacturing environment; they can inflate order quantities and contribute to excessive inventory levels.
- **Fixed Days Supply** This attribute, sometimes called “period of supply,” tells planning to generate planned orders to cover requirements for a certain number of working days. This will limit the total number of orders that are generated over time, but will result in overall higher inventory levels; consider using it for inexpensive items where the cost of ordering exceeds the carrying cost of the early inventory.

- **Safety Stock Method, Bucket Days, and Percent** The safety stock method MRP Planned % tells planning to dynamically calculate safety stock as the percentage of the average daily demand in the buckets you specify. Non-MRP Planned means that safety stock levels come from Oracle Inventory, either calculated by the Reload Safety Stocks concurrent program in Oracle Inventory, or manually entered on the Enter Item Safety Stock form.



NOTE

Non-MRP Planned does not imply that planning will ignore safety stock in its calculation; it simply identifies the source of the desired safety stock levels. Whether or not planning uses safety stock in its processing is controlled by the plan attribute Plan Safety Stock.

Mix-Max Quantity information is not used by MRP or ASCP. Cost information is used in ASCP optimization; Source information is one method of defining a supply chain; both of these topics are discussed in Chapter 11.

MRP/MPS Planning Attributes

The MRP/MPS Planning Attributes tab (see Figure 8-14) includes the fields that pertain specifically to MRP planning. The following fields are important:

- **Planning Method** The planning method determines the types of plans in which an item participates.
- **Forecast Control** This attribute controls forecast explosion and consumption, as described earlier.
- **Pegging** This determines the pegging information calculated by the plan. End assembly pegging maintains a “peg” or link to the end assembly responsible for the demand; Soft pegging maintains a link to the immediate source of the demand; Hard pegging is used for Project MRP. Various combinations of these values are allowed.
- **Exception Set** The exception set controls certain types of exceptions. To define exception sets on the Planning Exception Sets form, give each set a name and set the sensitivity and time period for six types of exceptions: Shortages, Excess inventory, Overpromised inventory, Repetitive variances, Resource over-utilization, and Resource under-utilization. You may have different exception sets for different classes of items; for example, your definition of “excess” inventory will probably vary for inexpensive, fast-moving items versus expensive, slow-movers. Note that if an item has no exception set, the exceptions listed here will not be reported for the item.

The screenshot shows the 'Organization Item (V1)' window for item 'MX344523' (Mother Board) under organization 'V1 Vision Operations'. The 'MPS/MRP Planning' tab is active, displaying the following settings:

- Planning Method: MRP planning
- Forecast Control: Consume and derive
- Exception Set: HiValHiVol
- Pegging: None
- Shrinkage Rate: (empty field)
- Acceptable Early Days: 2
- Round Order Quantities
- Repetitive Planning:
 - Overrun: (empty field) %
 - Acceptable Rate +: 0
 - Acceptable Rate -: 0
- MPS Planning:
 - Calculate ATP
 - Reduce MPS: (empty dropdown)
- Planning Time Fence: Total Lead Time (Days: empty field)
- Demand Time Fence: Cumulative Mfg. Lead Time (Days: empty field)
- Release Time Fence: User-Defined (Days: 5)

FIGURE 8-14. Item MRP/MPS planning attributes

- **Shrinkage Rate** The shrinkage is a decimal factor, usually small, that reflects an anticipated loss of the item during its manufacture or procurement; planning will inflate order quantities to account for this expected loss. Compare this with component yield, discussed later in this chapter.
- **Acceptable Early Days** This attribute filters out some “nervousness” in the plan. It suppresses “reschedule out” suggestions for existing orders if the order would be rescheduled by less than the number of acceptable early days. In other words, it defines how many days early an order can be before planning will suggest rescheduling it out.
- **Round Order Quantities** If enabled, this check box tells the planning process to round planned order quantities up to the next higher integer. Note that this affects planned orders only; it does not directly affect the calculation of dependent demand or of WIP requirements, which may be fractional due to yield factors.
- **Repetitive Planning attributes** Overrun % and Acceptable rate increase and decrease, as described earlier in the “Repetitive Planning” section.

- **Calculate ATP** This check box determines whether planning calculates and prints the available to promise on the Planning Detail Report; it does not affect the calculation of ATP for Order Management or Inventory inquiries.
- **Reduce MPS** This attribute lets you reduce master production schedule quantities based on the passage of time; use this only if you are not relieving the MPS through the creation of WIP jobs or purchase requisitions/orders (e.g., if you're using an external system for those functions).
- **Time Fences** You can specify a Planning Time Fence, which prevents new order creation or reschedule in suggestions; a Demand Time Fence, which you can use to ignore forecast demand either in planning or in loading an MDS; and a Release Time Fence, to define the window within which planning can automatically release orders from plans designated as Production. You define the length of each time fence either as a fixed number of work days or as the item's Total lead time, Cumulative total lead time, or Cumulative manufacturing lead time. Note that the use of the Planning and Demand Time Fences is further controlled by attributes of each individual plan.

Lead Times

Lead times in Oracle are divided into several elements. The relationship between these elements and the various types of lead time is shown in Figure 8-15.

- **Preprocessing** The time it takes to prepare an order for release; this might be the time it takes to prepare a shop packet for manufacturing or to prepare a purchase order. This is sometimes referred to as paperwork lead time.

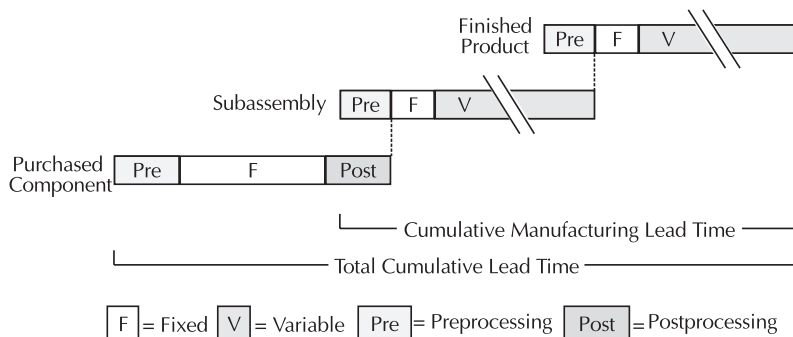


FIGURE 8-15. *Lead time elements and relationships*

- **Fixed** The portion of an item's lead time that does not vary with the order quantity. Setup time is a typical example; this is often a fixed amount of time, regardless of the size of the order. Planning uses fixed lead time to plan both manufacturing and purchase orders.
- **Variable** The portion of lead time that varies with the order quantity. Planning uses variable lead time to plan manufacturing orders only.
- **Processing** The time it takes for a supplier to fulfill your purchase order once released or the typical time it takes to complete a discrete job for the item. For manufactured items, this can be calculated from the fixed and variable lead times and represents the time it takes to manufacture the Lead time lot size quantity of the item; as such, it is an average. It is used to calculate the cumulative lead times. Planning does not use processing lead time at all for manufactured items; it uses the combined fixed and variable lead time. For purchased items, planning will use the processing lead time only if you do not specify a fixed lead time for the item.
- **Postprocessing** The time it takes to put away purchased material after it is received. Planning uses this for purchased material only; for manufactured items, you should include this in your routings if it is significant.

You can calculate fixed, variable, and processing lead times from your routing time, as described in Chapter 4. This process uses the lead time lot size, which you also specify on this tab.

The Lead Times tab also includes two types of cumulative lead time:

- **Cumulative Total** The "critical path" lead time, or the time it takes to buy the first piece of raw material, manufactures all the subassemblies and makes the item.
- **Cumulative Manufacturing** The cumulative lead time for manufacturing activities only; this is the cumulative total lead time, minus the purchasing lead time.

Cumulative lead times can be used to set the different planning time fences or the Infinite Supply Horizon in Oracle Inventory ATP rules. Using a cumulative lead time instead of a user-defined number of days means that these fences and horizons can shrink as you reduce lead times, without additional maintenance. Cumulative lead times are also helpful in reviewing proposed engineering changes; if you schedule a change inside an item's cumulative lead time, you may require expediting of material to satisfy the new demand. You can calculate cumulative lead times with the Rollup Cumulative Lead Times concurrent program after you have calculated or entered your item lead times.

Purchasing and WIP Attributes

The Purchasable and Build in WIP attributes control what types of orders you can release from the Planner Workbench. One or both of these must be enabled to allow you to release orders. The Make or Buy attribute and sourcing rules control the default type of order that is suggested, but you can override the suggestion with a different order type if the appropriate attribute is enabled.

BOM Information

The following component information from your bills of material is worth special mention in terms of planning:

- **Operation Seq** Identifies the routing operation at which the component is consumed. If you are planning by Operation Start Date, this information is used along with the Lead Time % from the item's routing to estimate when the component is needed. If your plan options specify planning by Order Start Date, this information is irrelevant for planning purposes; all material is planned to be available at the start of the job.
- **Planning Percent** Used in forecast explosion and in planning for Model, Option class, and Planning bills of material. The planning percent reflects the anticipated distribution of the component. Note that planning percents do not have to equal 100 percent. Certain optional components in a model will naturally total something other than 100 percent usage, and you can explicitly over- or underplan if you choose.
- **Yield** A decimal factor that represents the anticipated loss (or gain) of the component in manufacturing. Note that yield is expressed as a decimal factor and defaults to 1. A yield of .98, for example, represents that you expect to lose 2 percent of the component in the manufacturing process; planning will inflate the gross requirements of this component accordingly. Compare this with the item attribute Shrinkage; shrinkage represents an anticipated loss of the assembly in manufacturing or a purchased part during the procurement cycle. Planning will inflate the order quantity to account for shrinkage. It is expressed as a decimal quantity and defaults to 0. A shrinkage of .02, for example, indicates that you expect to reject 2 percent of the material during incoming inspection if it's a purchased part. These two factors, shrinkage and yield, are often confused. Though it is sometimes possible to manipulate the factors to give the same result, they are intended to model two very different events.
- **Material Control** For the most part, the material control attribute affects manufacturing execution, but two values have special significance in planning. A value of Supplier is intended to identify a part that is provided

by your supplier as part of outside processing; you are not responsible for procuring this part, so planning will not plan it. A value of Phantom invokes planning's phantom-part logic; it will "blow through" these parts and plan for their components directly. (Internally, planning actually does generate orders for the phantom, but these orders are visible on the Planner Workbench only if you query them explicitly.)

Planning Manager

The Planning Manager is a concurrent program (*not* a Concurrent Manager, despite the name); its function is to manage many of the "housekeeping" activities required in the planning environment. For example, the Planning Manager performs forecast consumption and master schedule relief. It also processes the forecast and master schedule interfaces, validating the data and populating the forecast and master schedules tables.

It is important that the Planning Manager be active. If it is not, you might see some unexpected results in planning. For example, you will not see discrete jobs or purchase orders for master schedule item in your plans; the rationale is these might not have been reviewed for master schedule relief, and if that relief did not occur, planning might overstate the master schedule.

The Planning Manager typically runs continually in a production environment, though the program itself executes on a regular interval; during the day it will execute and then schedule itself (with the same concurrent request ID) to run again. The next day, a new concurrent request ID will be assigned.

You can determine if the Planning Manager is active by examining the Active check box on the Planning Manager form, by using the View Requests form (the program should be running or pending), or by examining its log file for recent activity. You start the Planning Manager on the Launch Planning Manager form; the only parameter required is the Processing Interval. Oracle recommends starting with the default interval of 30 seconds.

Summary

This chapter covered the common features of single-organization MRP and multi-organization Advanced Supply Chain Planning. It described the planning process and basic forecasting, master scheduling, and capacity planning terminology and activities. This chapter described engineering's impact on planning and the unique characteristics of repetitive planning and kanban planning. It also covered basic setup steps, including the critical item and bill of material attributes that control the planning process.

The following chapters provide details on Oracle Demand Planning and the unique features of MRP and ASCP.